

AI and Machine Learning Tools at the Minnesota Supercomputing Institute


Benjamin Lynch, Ph.D.
Associate Director for Scientific
Computing, MSI



Spock: “Computer, digest log recordings for past five solar minutes. Correlate hypotheses. Compare with life forms register. Question. Could such an entity within discussed limits exist in this galaxy?”



Computer: “The Drella of Alpha Carinae Five derives nourishment from the emotion of love. There is sufficient precedent for existence of creature, nature unknown, which could exist on emotion of fear.”





Gwyn: Alexa, play some Justin Timberlake.

Alexa: Playing songs by Justin Timberlake.



 "Ah, yeah, ah, yeah
I got this feelin' inside my bones
It goes electric..." 



Ben: Alexa, play “En Mis Noches” by Balún.

Alexa: I’m sorry, I can’t find the song “In the Snow Chess” by Berlin

Deep Neural Networks




Outline

- MSI Hardware & Software
- MSI Access
- Storage at MSI
 - Home directories
 - Local scratch
 - Tier 2 Storage (S3)
- AI/ML Tools at MSI
 - Running ML tools through Open OnDemand
 - PyTorch Example

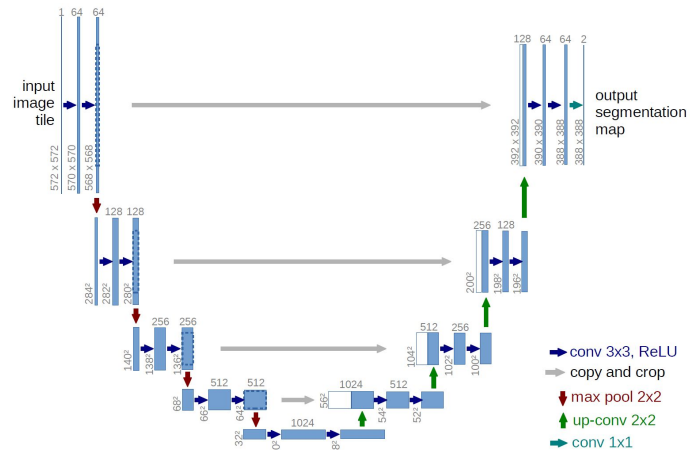

Minnesota Supercomputing Institute

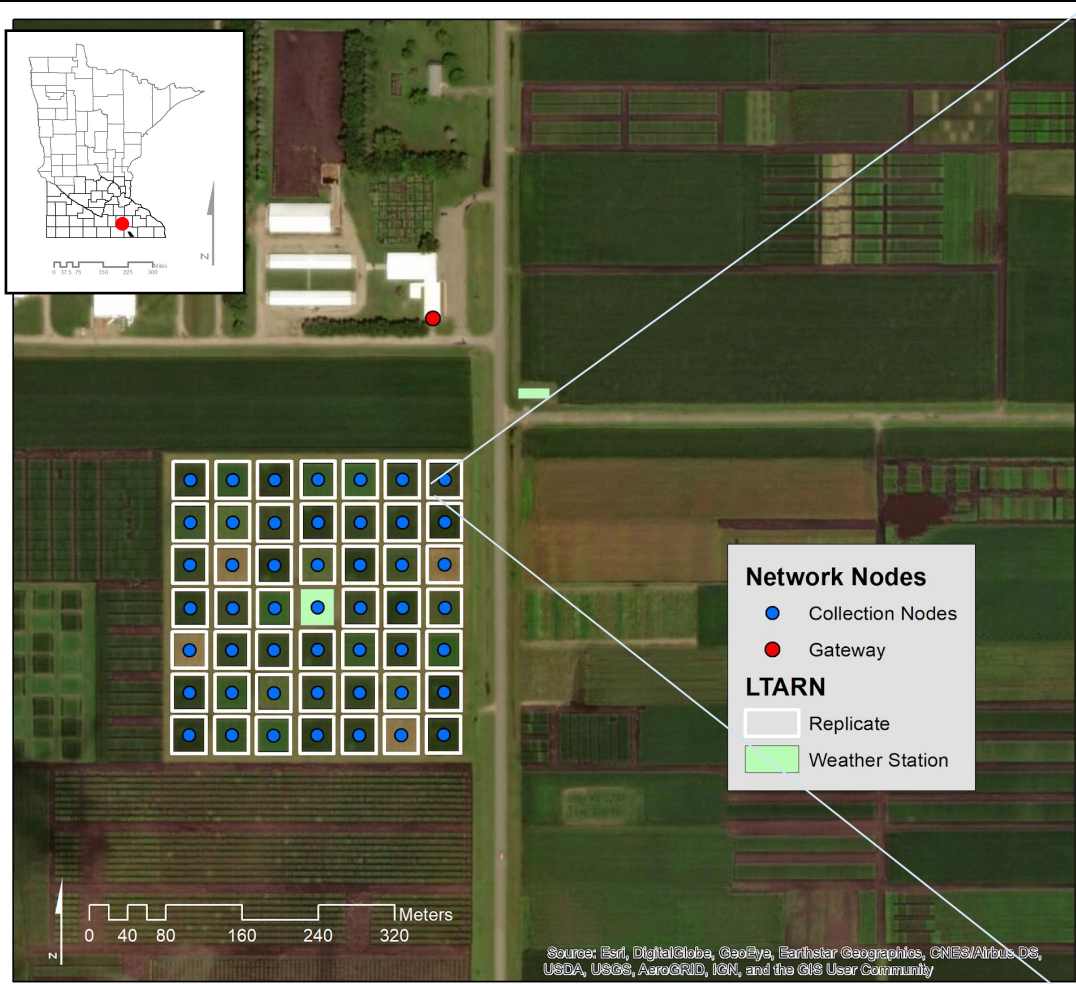
- Hardware
- Software
- Consulting
- 56 staff and growing

MSI Hardware



AI/ML Applications






Chlorophyll-a

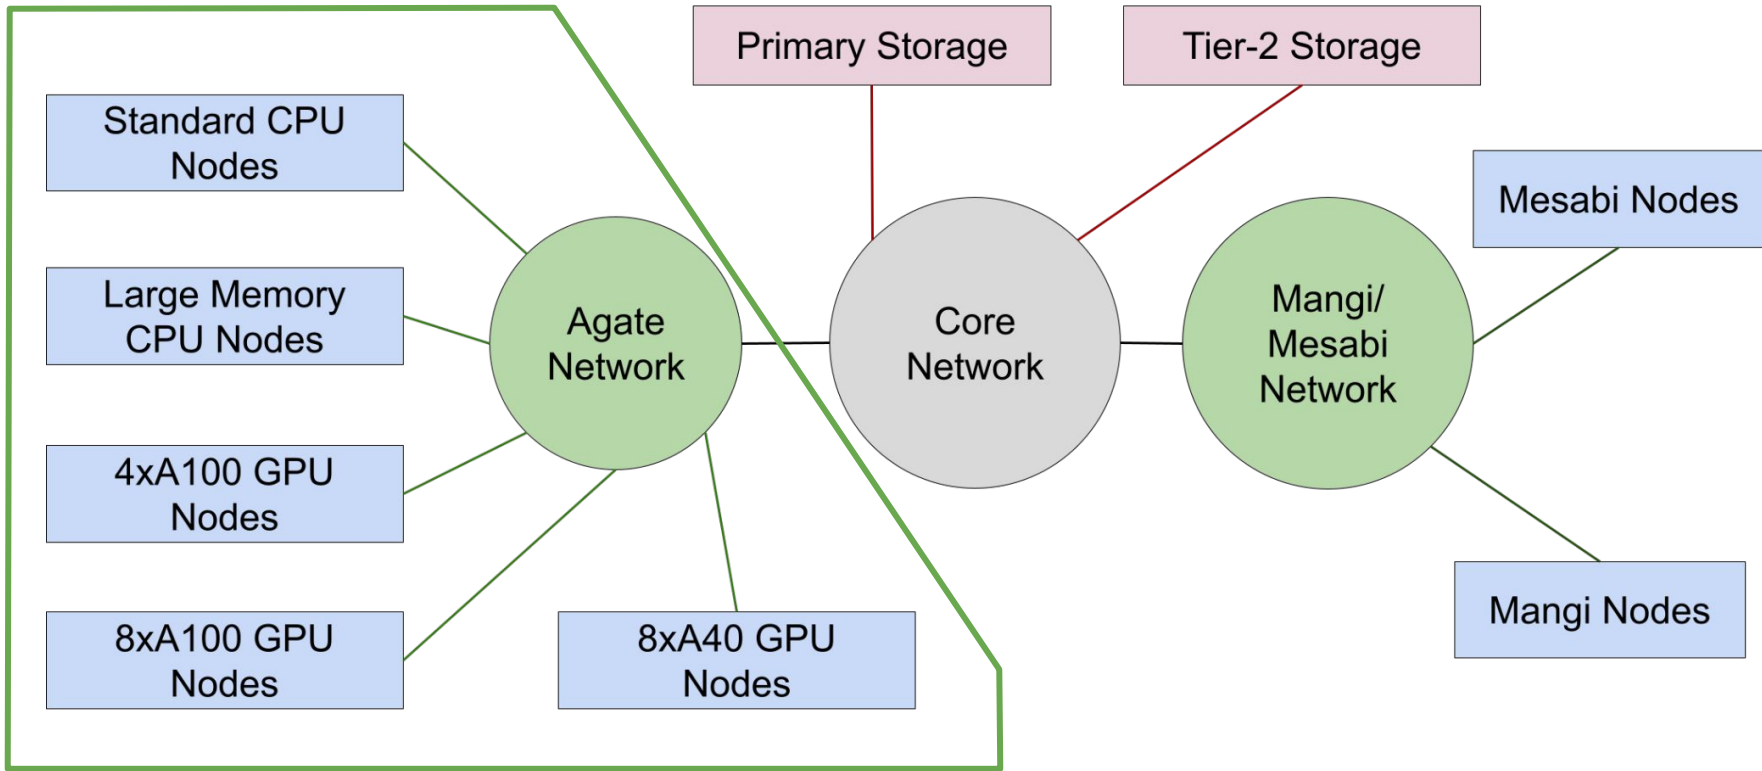
CARLSON'S
TROPIC STATE INDEX (TSI)

Chlorophyll-a TSI log scale

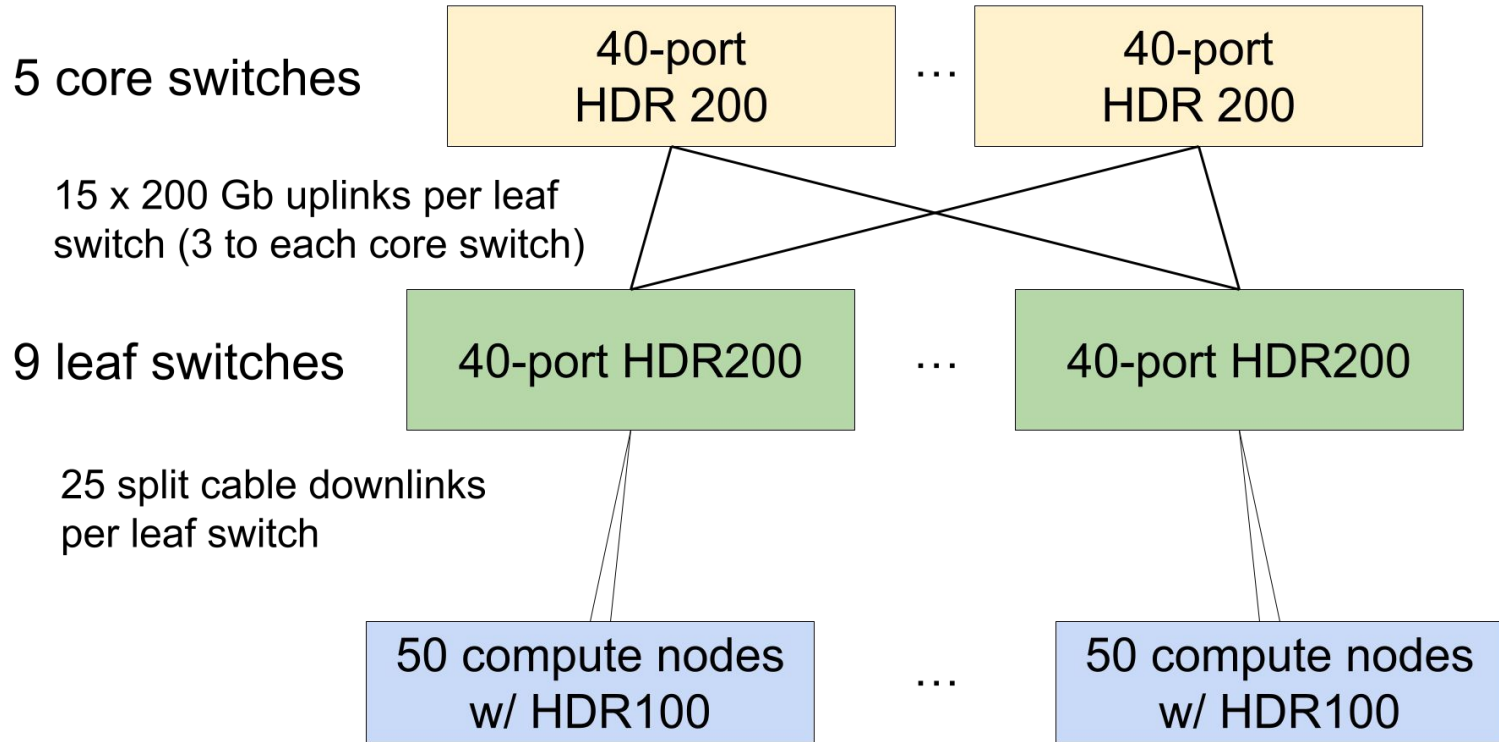
0	50	100
0.044	7.22	1181 [mg/m ³]



HPC Resources

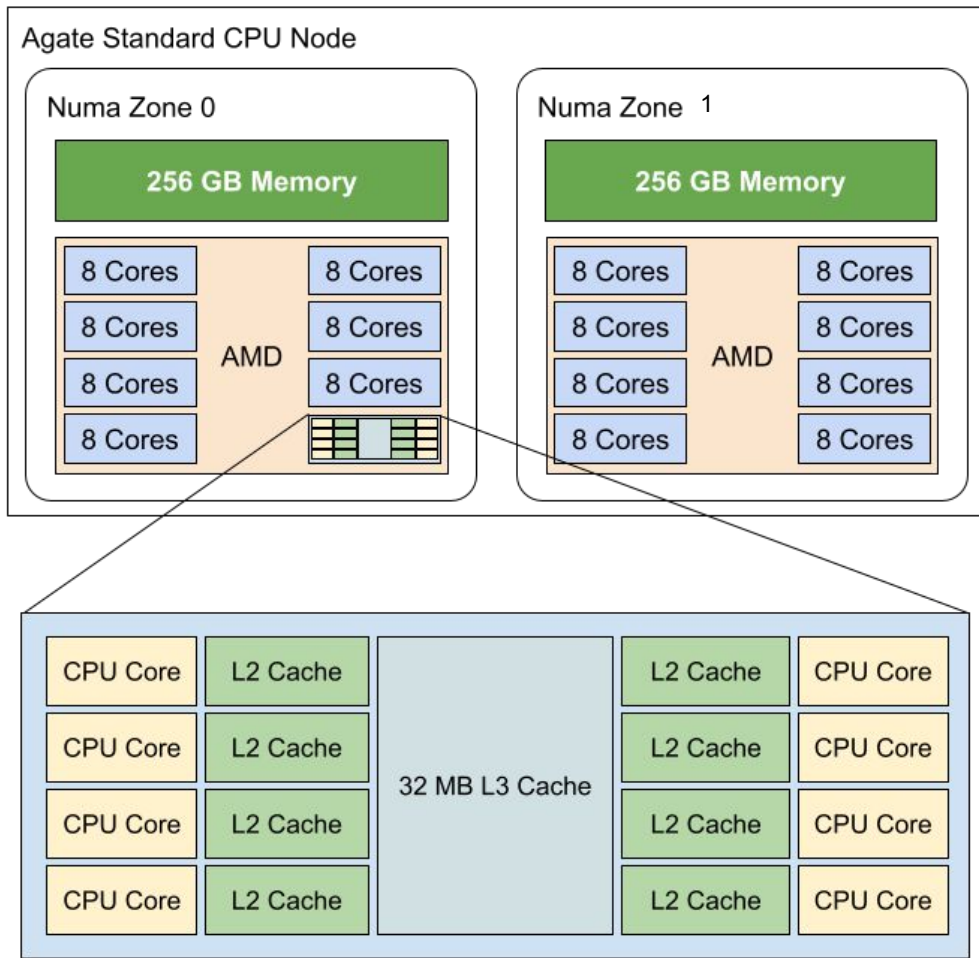


Agate Infiniband Topology



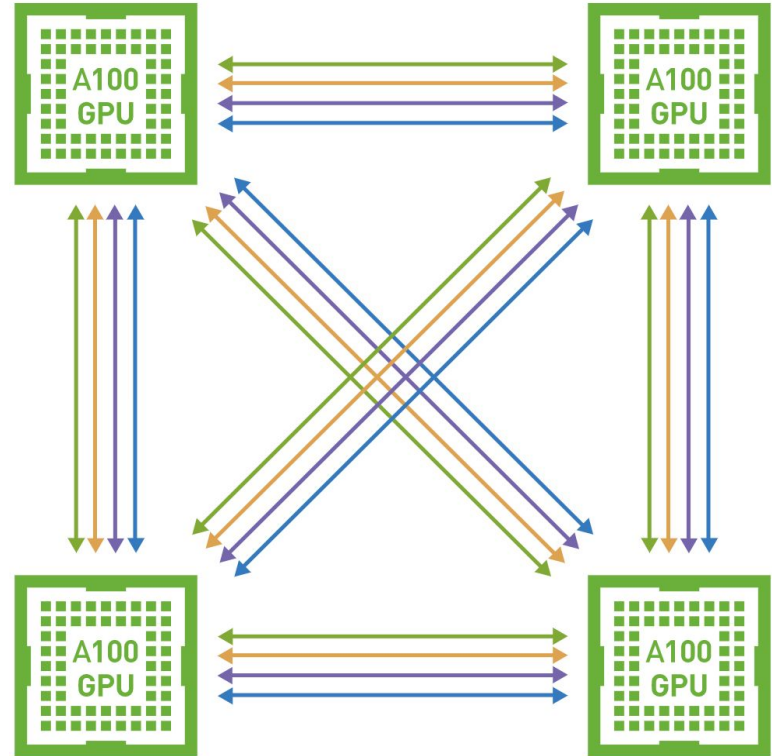
Standard Agate CPU node

- 128 CPU cores
- 512 GB of memory
- 850 GB local SSD

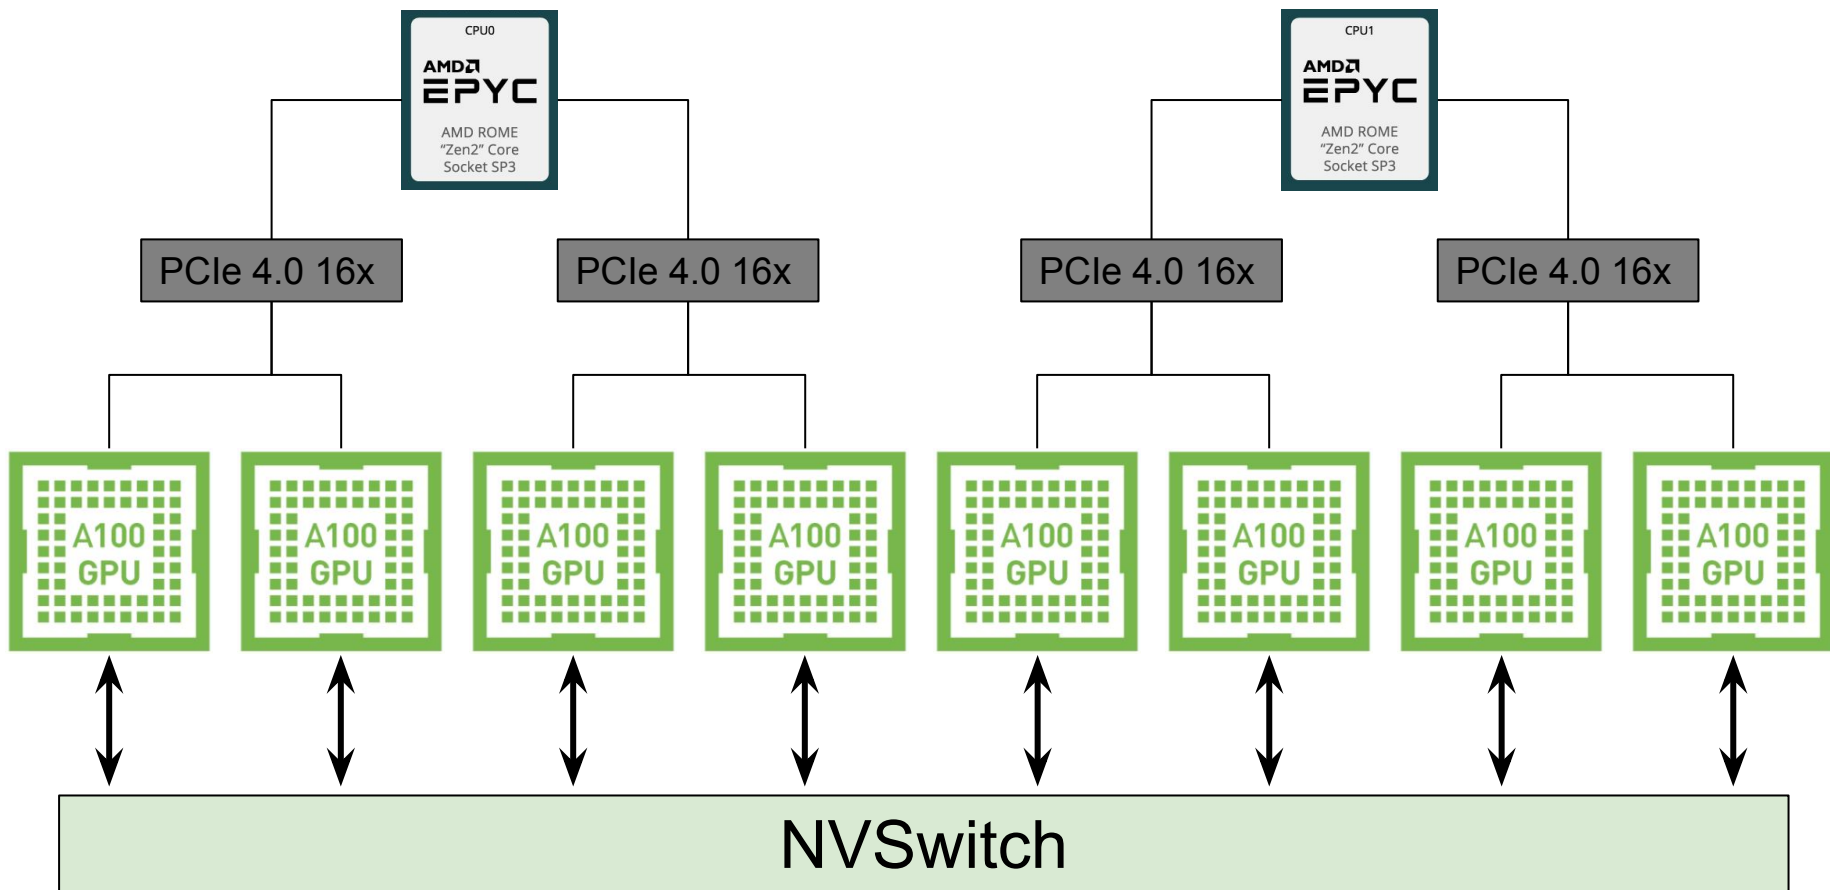


A100-4

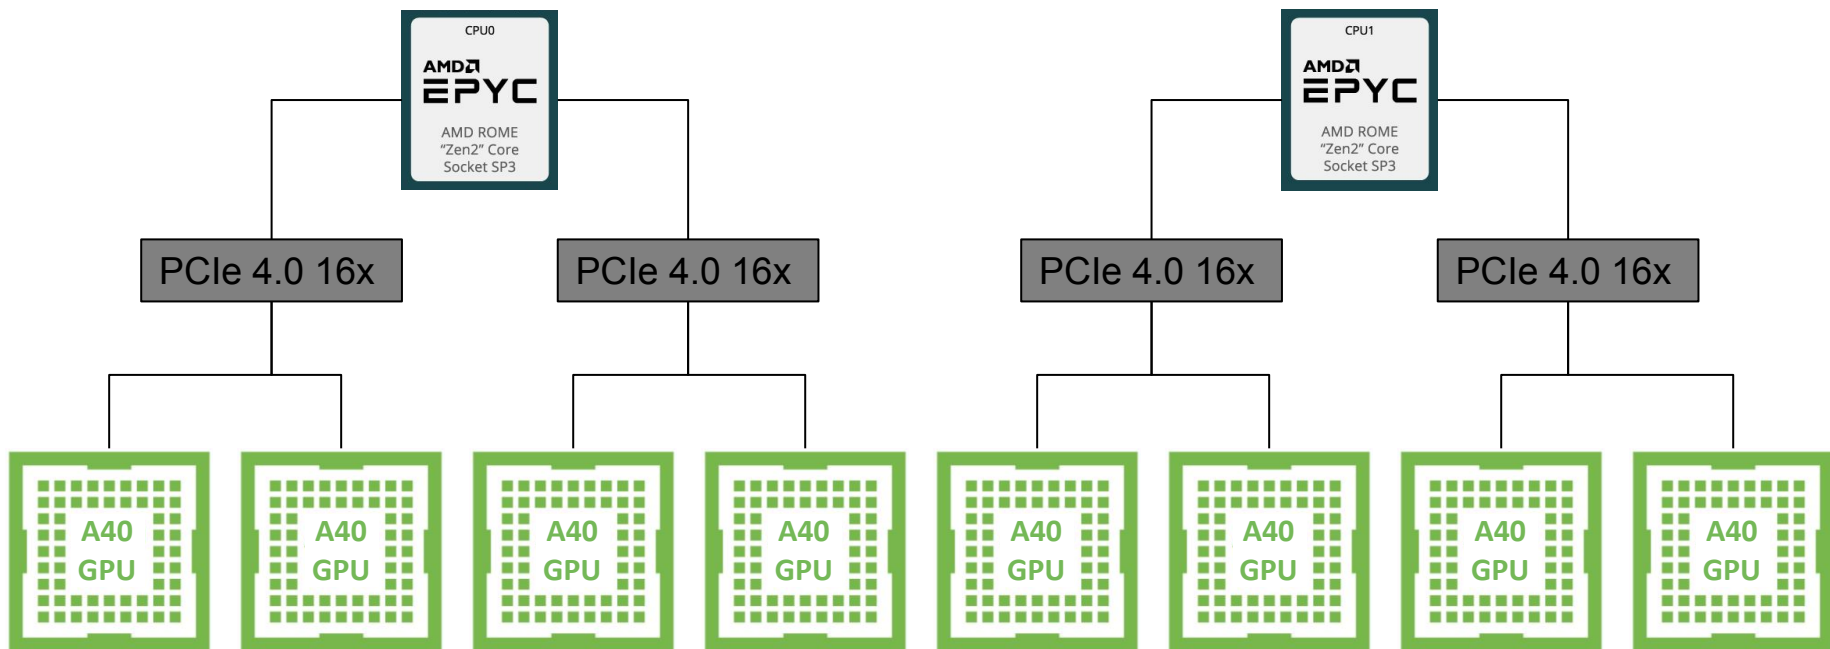
- 512 GB memory
- 64 CPU Cores
- 4 A100 GPUs
 - 3rd Generation NVLink
 - 600 GB/s bandwidth per GPU



A100-8



A40-8



MSI Storage

Storage Location	Located at	Appropriate Size	Notes
Local Scratch	/tmp	10s of GB	Ephemeral
Home directory	~/	100 GB	Snapshots available
Tier 2 Storage	s3.msi.umn.edu	TBs	Bulk storage, no backup

Connect to MSI

<https://ood.msi.umn.edu>

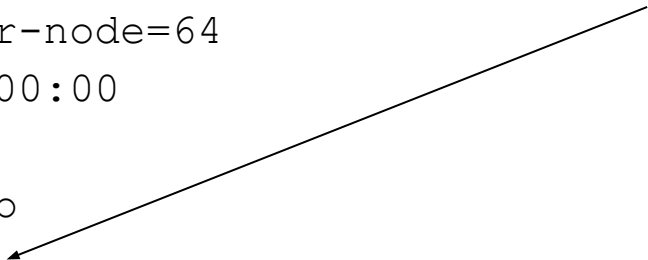
OR

```
ssh agate.msi.umn.edu
```

Slurm Script

```
#!/bin/bash -l
#
#SBATCH --mail-type=NONE
#SBATCH -N 1
#SBATCH --tasks-per-node=64
#SBATCH --time=24:00:00
#SBATCH --mem=400G
#SBATCH --tmp 400gb
#SBATCH -p a100-4
#SBATCH --gres=gpu:a100:4
```

Specify the 4-way GPU nodes if you want to ensure your job lands on this node type.

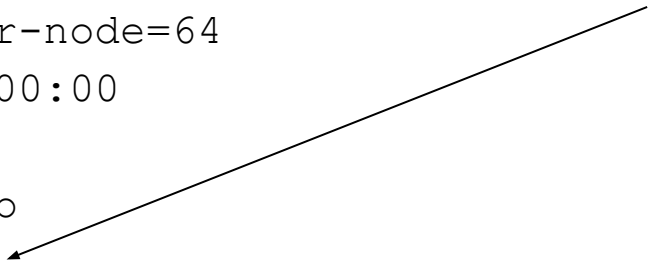


```
do_something
```

Slurm Script

```
#!/bin/bash -l
#
#SBATCH --mail-type=NONE
#SBATCH -N 1
#SBATCH --tasks-per-node=64
#SBATCH --time=24:00:00
#SBATCH --mem=400G
#SBATCH --tmp 400gb
#SBATCH -p msigpu
#SBATCH --gres=gpu:a100:4
```

Specify the **msigpu** partition if you don't care of the job lands on a **a100-4** or an **a100-8** node.



```
do_something
```

Open OnDemand

<https://ood.msi.umn.edu>

- MSI has recently deployed a new portal for researchers to access MSI systems.
- From this interface, researchers can get shell access to the HPC login nodes

Documentation

<https://www.msi.umn.edu/content/pytorch-open-ondemand-ood>

<https://www.msi.umn.edu/content/custom-jupyter-notebooks-using-singularity>

<https://www.msi.umn.edu/support/faq/how-do-i-use-second-tier-storage-command-line>

Google CoLab Example:

<https://colab.research.google.com/drive/1W9C2YDktfjyTAwhtZxl3BqCU192lCYjw#scrollTo=tTBsQ7Vzu63>

<https://www.msi.umn.edu/quick-start-guides>

<https://msi.umn.edu/software>

This Presentation

<http://z.umn.edu/csci5527>

Basic PyTorch Example

wget <https://public.s3.msi.umn.edu/tutorials/csci5527/PyTorch-MNIST.ipynb>

Job

We are looking for PhD or MS students for Fall term to work on developing a 3D medical imaging visualization project. The goal is to develop a graphical user interface that allows a user to interact with multiple 3D volumetric images acquired using MRI, and prescribe a 3D region of interest within a tumor. Depending on skill set the program may be a standalone application or a web-based app. The initial appointment may potentially be extended for the Spring term and beyond. Please note that you must be physically in the United States to be eligible for this position.

Required qualifications:

Experience with VTK or other 3D visualization software toolkits

Proficiency in program GUIs and algorithms in either Matlab, python/Qt, or javascript/HTML5

Preferred qualifications:

Experience with medical image analysis

Please contact Malgorzata Marjanska (gosia@umn.edu) if you are interested.