

Welcome to the AiMOS Supercomputer



Chris Carothers, CCI Director
Dan Weeks, Lead HPC Developer
Jay McGlothlin, Manager, Research Computing
Derek Fox, System Admin II
Rensselaer Polytechnic Institute



Thank You!! - AiMOS Was A Team Effort



Welcome AiMOS! – The AI Hardware Center’s Supercomputer

AiMOS: Artificial Intelligence Multiprocessing Optimized System

- Will serve as the supercomputer testbed for the **multi-billion AI Hardware Center** initiative in New York State

Name has several additional meanings:

- “Amos Eaton” – first senior professor at Rensselaer
- “AI” – supporting research & educational priorities in AI & ML
- “MOS” - New metal-oxide semiconductors for AI/ML.

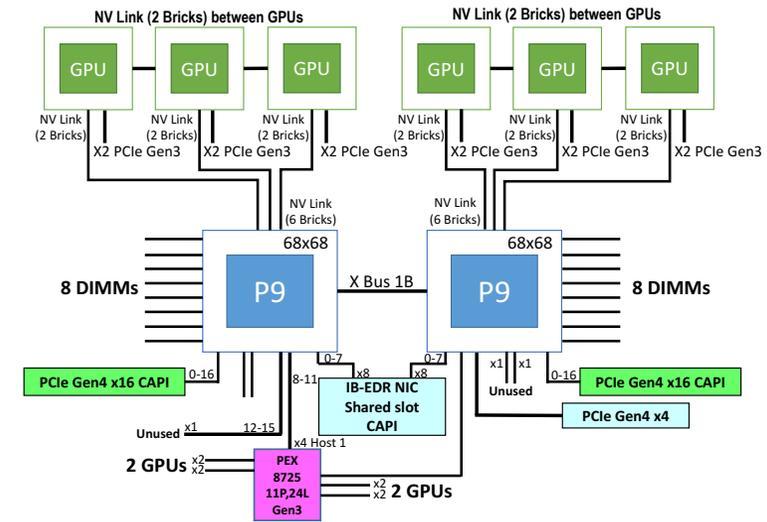
The current configuration for AiMOS is:

- 252 AC922 compute nodes configured in a 14 rack system.
- Each Compute Node has:
 - **2x 20 core POWER9 processors (Summit’s variant)**
 - **6x NVIDIA Volta V100 GPUs w/ 32 GB of HBM each**
 - 512 GB of DRAM
 - 1.6 TB NVMe SSD storage
 - Dual, 100 Gb/sec Mellanox IB links
- Additional, 16 computes node w/ 4 GPUs are part of NSF/MRI award.
- Fat tree network has MOR IB switches connected to 2 Director switches
- **All systems use 480V power (1 PDU per rack) & water cooled!**
- ~11 PB of IBM ESS Storage @ ~80 GB/sec read



HPL Performance

- Performance is: **8.045 PF & 15.771 GF/watt @ 250 compute nodes**
- Much harder than we initially thought to obtain
- HPL benchmark binary obtained from IBM & NVIDIA
- #24 Top500
- #3 Green500



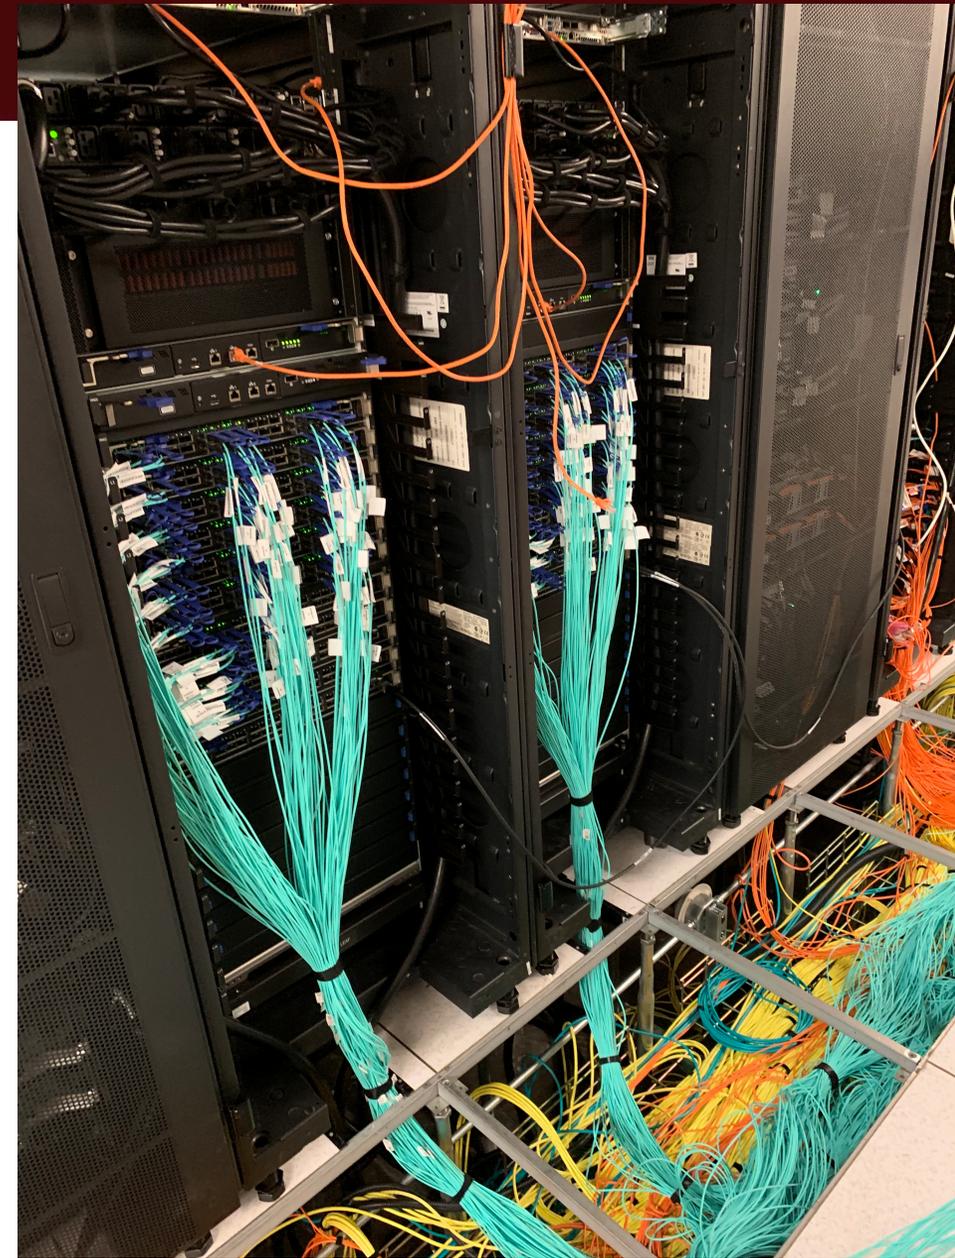
HPL 250 Node Run on AiMOS



Power Measurement

What about the Mellanox Director switches ?

- We used the Mellanox switch web interface to capture the power (in KWs)
- Observed that network traffic never above 5%, zero congestion and switches never went above 3.5KW each or 7KW total



Putting HPL core measurements together...

- Total compute node, MOR switch energy usage for HPL core:
 - 648 KWH (372 KWH + 276 KWH)
- Total HPL core time:
 - 4636 secs or 1.288 hours
- Avg. KW for compute nodes and MOR switches:
 - $648 \text{ KWH} / 1.288 \text{ H} = 503 \text{ KW}$
 - **System is designed for ~700 KW**
- Avg. KW for Director switches:
 - 7 KW
- Total KW = 510
- Overall power efficiency:
 - **8,045 TF/510 KW = 15.8 GF/watt**

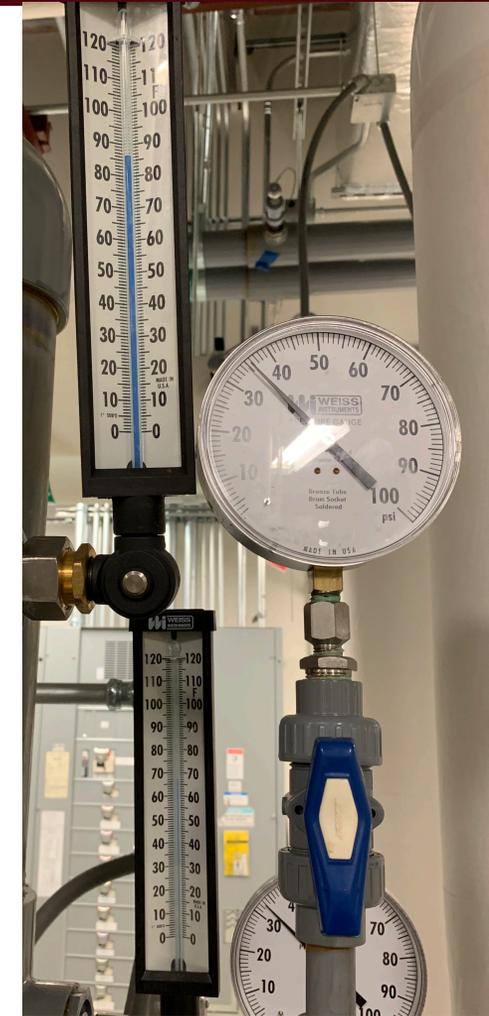


Big Challenge: GPU Throttling

- Started with overclocking GPUs & CPUs
- Initial water cooling setup used 75 F degree water
- Observed on many runs that GPUs were power throttling.
 - E.g., trying to exceed more than 300 watts
 - Resulted in power supplies and whole system pulsing thru the HPL run.
 - Thought it was a few bad GPUs but they all passed DGEMM and other tests.
 - Almost to the point of giving up on 8.0 PF
- Fix: lower water temp to 65F @ zero cost (free cooling) and issue went away.
- Take away: Cold Water Matters!!



75 F in, 90F out



65 F in, 84 F out

Thank You!



Rensselaer