

MEET THE TEAM



Ryan Scherbarth
Senior, CS & Math
Focus: Benchmarks
Working concurrent positions at Sandia National laboratories on the production HPC team, and at UNM at the Center for Advanced Research Computing & Teaching Asst. for CS491: HPC.



Sumaya Houssini Mohamed
Senior, CS & Business
Focus: NAMD
Currently working at the Center for Advanced Research Computing, and previously involved in multiple competitions with the UNM HPC team including ISC24 competition.



Xavier Barr
Senior, CS & Math
Focus: Benchmarks
Working as a Robotics Research Faculty Assistant in the Moses Biological and Computation Lab, and previously competed with UNM HPC in the Winter Classic 2024.



Saskia Trefethen
Junior, CS & Business
Focus: ICON
Intern at COSMIAC working on embedded software for flight vehicles, and previously completed multiple HPC related courses at UNM including parallel computing and HPC.




Reece Ransby
Senior, CS & Math
Focus: NAMD
Working at Sandia National Laboratories, participated also in the Winter Classic competition with the UNM HPC team in the spring.



Molly Palko
Senior, CS & Math
Focus: ICON
Joined the UNM HPC team in Spring 2024, and has completed related courses at UNM including HPC, and works at AFRL in HPM modeling.

OPTIMIZATION STRATEGY


Our optimization strategy starts with our hardware selection. Benchmarking per-watt performance on our accelerators prior to the competition helps us determine the optimal hardware configuration given the competition constraints.



Graph showing efficiency increases as power is throttled

Applications are distributed across team members who each have a primary focus. We are meeting weekly where we all sync on application progress and work together on high priority issues.

We have the ability to work using UNM CARC resources before the cluster arrives and in the time between shipping and competition start.



H100 Accelerator is most efficient at 50% power throttle

We plan to utilize Nvidia optimized versions of HPL and MLPerf as a starting point, as well as implementing what we have learned about these benchmarks in our previous HPC competitions.

Our team has experience with HPL through multiple competitions, MLPerf in SC23, and ICON in ISC24. We will be utilizing this experience and saved scripts / notes.

HARDWARE & SOFTWARE


Our three node cluster is centered around GPU compute with 4x H100 NVL accelerators per node. We chose the NVL variant for a much better balance between performance and power consumption than the SXM;

- 12% lower peak performance (F64)
- 18% more memory
- 43% lower peak power draw

Each node has 2x Intel Xeon 8580 CPUs. With hyper-threading enabled there are 240 cores per node, and slightly over 2TB DDR5-5600 RAM. We are using warewulf for provisioning the Rocky 9 OS image, and slurm for scheduling.

Our cluster is well equipped for the applications this year, with GPU compute as the backbone. HPL and MLPerf benchmarks are GPU bound, ICON has a recent GPU-accelerated release, and NAMD likewise will have a significant demand for GPU & CPU power.

Each GPU is connected with a 1800 GB/s NVLink connection, and each node has 4x Nvidia ConnectX-7 Network Interface Cards.



HPC AT UNM

The University of New Mexico, a tier one research institution located in Albuquerque, New Mexico has a strong interest in the continued advancement of high-performance computing. UNM strives to provide students with the opportunities and resources to start their careers in HPC through workforce training, student activities, and jobs.

The Roadrunners HPC team founded in January 2023 under Dr. Matthew Fricke with the ambition to foster students' interest in HPC while helping them develop applicable skills. The Roadrunners team showcases UNM's diversity with over half of our members identifying as a part of traditionally underrepresented groups in STEM fields.



FACULTY ADVISOR
Dr. Matthew Fricke
TEAM COACHES
Nick Bacon, Alex Knigge

