



Computational Facilities

The Advanced Research Computing Center (ARCC) located at the University of Wyoming (UW) is the university's primary research computing facility. All ARCC operations promote the growth of UW's research and education activities. Resources provided by the department include high performance computing (HPC), large research data storage, and user training and consultation. Specialized ARCC staff members perform on-site administration, maintenance and support on all ARCC hosted systems and research support resources, including UW's primary HPC cluster, Beartooth. This is a heterogeneous condominium cluster open to all facets of research. ARCC also supports several specialty clusters and research services used by UW research organizations. HPC hardware is housed in liquid-cooled enclosures in UW IT's 6,000 sqft data center which includes highly redundant infrastructure for power, cooling, and security. Hardware is interconnected using Mellanox InfiniBand and high-speed ethernet.

Beartooth HPC System Configuration

Beartooth is a x86_64 based HPC running RHEL 8. Job scheduling configuration is based on the condo model which encourages researchers to purchase nodes for use in our HPC environment granting priority and predictable access. ARCC leverages fair-share mechanisms to distribute computational workloads between varied research projects as well as weighted scheduling parameters (job size, age, etc.) to maximize efficient cluster utilization. Beartooth is backed by VAST performance data storage platform hosting over 3 PB of high-performance storage.

Beartooth (currently at over 375 nodes) hardware is heterogenous, comprised of several specialty partitions to address a wide variety of scientific pipelines and workloads. These include mass memory (4TB RAM) nodes, and a variety of GPU nodes. The system utilizes SLURM workload manager and LMOD environment modules to provide a robust and flexible user experience. Beartooth supports a wide range of compilers (GNU, Intel oneAPI, and NVIDIA HPC SDK) as well as containerization frameworks.

Beartooth hardware specifications are summarized below:

Beartooth	#Nodes	Cores	RAM (GB)	GPU (mixed)	GPUs / Node
Hardware Specs	384	15468	84504	52	Up to 8

Facilities Expansion: Thunderer

Thunderer, an upgrade to ARCC's primary HPC environment, is made possible by a generous \$5 million appropriation from the State of Wyoming. Hardware specifications for the addition are listed in the following table:

Allocations	#Nodes	Cores / Node	RAM	GPU	GPUs / Node	Processor	TensorCores/GPU & CUDACores/GPU
CPU	25	96 CPU Cores per Node	1024GB/Node	N/A	N/A	2x 48-Core/96-Thread 4th Gen AMD EPYC 9454	N/A
A30 GPU	8		24GB/GPU	Nvidia A30	8		224 TC/GPU 3804 FP32 CUDA/GPU
L40S GPU	5		48GB/GPU	Nvidia L40S	8		568 TC/GPU
H100 GPU	6		80GB/GPU	Nvidia SXM5 H100	8		528 TC/GPU 16896 FP32 CUDA/GPU

High Performance Data Storage

Our facility houses two main research storage resources in addition to housing data for direct use within Beartooth. Both are available to all researchers at UW for both short- and long-term archive data storage. The first is named Alcova, a peta-scale capable system. Both Beartooth and Alcova file storage use the VAST data platform, a highly scalable, in-line, block-based data storage service, built for performance. Connected to the UWyo network, Alcova offers blazing-fast transfer speeds starting at 100 GB/s. Alcova also facilitates the publication of datasets curated by UW Libraries. Pathfinder, ARCC's other primary research storage resource, uses the S3 protocol on object storage through Ceph to provide a low-cost storage option with added functionality to allow access to wide scale of large datasets for web-based applications. All ARCC hosted data storage services support GridFTP via Globus data transfer servers connected to the UW Science DMZ (100 Gbps Internet2 link). Alcova and Beartooth storage also allow for SMB/CIFS and NFS making it easy for researchers to transfer data from their daily work environment.

In addition to computational resources, our department assists researchers by offering guidance, consultation and training for any users - both new and experienced - who may seek help incorporating HPC technologies in their research pipeline. Our team works in collaboration with the UW Libraries Digital Scholarship Center to host training towards the growth in understanding and use of HPC across UW and throughout our state.