FACILITIES and EQUIPMENT
for RESEARCH COMPUTING at the UNIVERSITY OF WYOMING

FACILITIES

Advanced Research Computing Center

The Advanced Research Computing Center (ARCC) at the University of Wyoming (UW) provides computational hardware and software support for research computing, including high performance computing (HPC), as well as large data storage, training and consultation to further UW's research and educational activities. ARCC provides systems administration for the campus HPC condominium cluster Beartooth, which is freely open to all facets of research. In addition, ARCC supports numerous specialty systems and a three-tier data storage system, to accommodate a variety of unique research needs. ARCC develops, hosts and delivers training courses, workshops, research consulting and internship opportunities in compute-intensive and big data projects for graduate and undergraduate students in diverse fields, including computer science, physics, astronomy, mathematics, biological sciences, English, geology and other areas. The work is conducted by six specialists in hardware architecture, system administration, software and user support, co-led by the UW CTO for Research and the ARCC Director.

IT Data Center: A state-of-the-art, 6,000 square foot facility hosts several hundred production computing systems that provide critical information technology (IT) services required to operate the UW, as well as the ARCC HPC equipment. The Data Center is a multi-zoned space. This design offers traditional raised floor cooling for low to medium density requirements that are common in utilitarian, mixed solution type facilities. Within the same envelope, the center also provides a high ampacity electrical distribution system coupled with closed loop cooling. This meets the demands of high-density deployments needed for intense research, supercomputers and cluster computing. The Center has approximately 150 cabinets for servers and networking gear. Virtualization is heavily utilized to improve capacity, density and power/cooling efficiency.

The electrical distribution system uses utility power as the primary source while providing backup power via an emergency 1.2-MegaWatt Generator. Automatic transfer switches (ATS Units) and parallel redundant uninterruptible power systems (UPS units) provide switching and ride-through during power transitions between the primary and secondary sources. In a steady state, the UPS System operates at ~99% efficiency. On the data center floor, power is delivered via 208V branch circuits. Each cabinet is equipped with 208V 30Amp single phase or 60Amp 3Phase power strips, which only provide specialized C13/C14 and C19/C20 receptacles.

The primary means of heat evacuation for the IT Data Center is an onsite specialized "free cooling" chiller with a secondary/backup campus chilled water loop. The campus cooling loop also offers cold water storage and outside air/evaporative cooling to improve overall efficiency.

The Data Center is a secured space that includes perimeter controls, hour restrictions, dual factor biometrics, video surveillance/recording, and traditional key- and combination-based locking. Numerous network- and system-based security measures are also deployed.

The UW Data Center utilizes a full complement of tools for management, monitoring, automated controls and state-change alerting. These tools are essential elements of the Data Center and enhance problem recognition, capacity planning and performance management.

Networking

The University of Wyoming is a founding member of the Front Range GigaPop (FRGP) connecting education and research across Colorado and Wyoming at 100s of Gigabit of bandwidth. Since 1999, FRGP has advanced the research and educational goals of government, nonprofit, research, and educational participants in the region by establishing and maintaining a unique multi-state network. The FRGP is owned and controlled by the FRGP participant research and education community.

In addition, UW is a founding member of the BiState Optical Network (BiSON), a dense wave division based optical ring connecting many of the FRGP entities together and to Internet2 and the world. In 2016 UW became the first university in the Rockies to connect to I2 with a 100G circuit via BiSON funded by a CC*NRI grant from NSF.

The University of Wyoming Science Network (UWSN) is the campus implementation of the Science DMZ principle. A Science DMZ ("science demilitarized zone") is a portion of a larger network that has been configured and optimized for high-volume bulk data transfer, remote experiment control, and data visualization for high-performance science applications. The guiding principle behind UWSN is that the campus research community should be able to optimize their access to other research entities, data stores, and computing resources specific to the needs of their individual projects. To accomplish this, they must overcome barriers of bandwidth constraint, high-latency, or the restrictions of an active security perimeter and be provisioned with the most efficient network possible. In order to achieve the maximum speed and throughput possible, the UW implementation of the science DMZ model avoids the regular campus exit architecture and is neither restricted nor protected by the campus firewalls.

IT Development Laboratory: A large, air-conditioned room of over 600 square feet, the development lab houses a row of 7 compute racks with space for additional rows, as well as the dedicated equipment staging area, appropriate power delivery mechanism and office space. The lab is set up for the purpose of developing new computer system architectures and testing them before deployment in the Data Center.

Animal facilities: not applicable

Clinical facilities: not applicable

ARCC office computer resources: Every Staff member at ARCC, including student interns and post-doctoral associates, is assigned modern office equipment, such as a computer, flat screen monitor, ergonomic keyboards and mice, as well as access to shared network printers, scanners and copier machines.

ARCC Library: A 315 square foot room is dedicated as a study and meeting area. It houses several shelves filled with books on relevant computing, scientific and leadership subjects, freely accessible to each member of ARCC, including students and post-docs. It also contains enough office desks and seating to conduct large meetings (up to a dozen participants) and remote teleconferencing via a computer hooked up to large screen TV and audio-visual equipment.

ARCC Staff Office space: The Advanced Research Computing Center occupies office space on two floors of the UW IT building, with 5 offices dedicated to student interns and post-doctoral researchers, and 5 offices dedicated to system architects, administrators, software and user support specialists. Each office is over 140 square feet in size, and equipped with variable-height

desks, ergonomic office chairs, wardrobes, shelving and other storage space. Remote working capabilities are also enabled.

ARCC leadership office space: The ARCC Director occupies a 256 square foot office on the third floor of the IT building, in proximity with the student offices and the IT Leadership. The room is set up with both office space for the Director and meeting space for up to 10 participants, 2 whiteboards, shelving and a wardrobe. The Deputy Director occupies a 144 square foot office on the second floor of the IT building, in proximity to the ARCC Infrastructure staff and the Data Center. The office is equipped with electric variable-height desk, office chairs, a white board, shelving and a small meeting space for up to 3 participants.

MAJOR EQUIPMENT

ARCC HPC cluster

Beartooth is an Intel x86_64 based HPC system consisting of over 550 nodes which include Sandy Bridge/Ivy Bridge, Haswell, Broadwell, Knights Landing, Cascade Lake, Skylake and Ice Lake architectures, as well as AMD Epyc, all utilizing a Mellanox InfiniBand FDR&EDR interconnect. Varying scientific workloads can benefit by having heterogeneous computing components. Beartooth hosts specialty nodes to address the many different computational requirements, including (1) huge memory (4TB RAM) nodes, (2) GPU nodes with NVIDIA K80 and A30 GPUs and a range of memory, from 256 GB/node to 1 TB/node, as well as (3) multiple Nvidia DGX nodes with 8 Tesla V100 GPUs each. Beartooth utilizes the SLURM workload manager and Lmod environment modules to provide a robust and flexible HPC user experience. The system supports a wide range of compilers (GNU, Intel, LLVM, and PGI), Allinea tools like DDT and MAP for advanced debugging and profiling, as well as containerization frameworks like Singularity. The system job scheduling configuration is based on a condominium model that encourages researchers to purchase compute nodes in the system granting additional priority and predictable access times. Beartooth leverages fairshare mechanisms to fairly distribute compute workloads between research projects, as well as weighted scheduling parameters (job size, age, etc.) to utilize the cluster as efficiently as possible. Hundreds of scientific and programming software packages are maintained on the cluster, at the request of UW researchers.

ARCC Storage

Beartooth is backed by an IBM Elastic Storage System 3000 & 5000 pair which run IBM Spectrum Scale (formally GPFS) parallel filesystem hosting approximately 1.2 PB of high-performance storage, designed for researchers to perform computing against it. This storage system supports GridFTP via Globus data transfer servers that are connected to the UW Science DMZ (100 Gbps Internet2 link) at 40 Gbps. Beartooth also provides SMB/CIFS and NFS for researchers to transfer data from their daily work environment. In addition to this fast, performance grade storage, ARCC also offers Alcova (over 1 PB) – highly collaborative space geared toward project-oriented data storage for use by UWYO research groups and their collaborators. This reliable storage tier provides easy access from Windows and Mac and data sharing capabilities via Globus Online and SMB/CIFS. Finally, Pathfinder (over 1 PB), named after one of Wyoming's reservoirs on the North Platte River, is a low-cost, expandable storage solution that enables a cloud-like presence for research data hosted by ARCC. Its core functionality is hosting onsite backups as well as enabling data sharing and collaboration. Pathfinder uses the Simple Storage Service (S3) protocol originally developed by Amazon. S3 works on object storage through a service called Ceph, provided by Red Hat Enterprise Linux. These characteristics make Pathfinder uniquely suited to providing a cloud-like data repository.

WildIris, the INBRE virtual cluster

WildIris is a specialty cluster designed to enable outreach for INBRE researchers and Wyoming community college students. It can be securely accessed anywhere, anytime using SSH connectivity with two-factor authentication. It has a hybrid architecture, consisting of 4 virtual nodes (16 cores, 64 GB RAM) and 4 physical nodes (48 cores, 512 GB RAM), in addition to the standard login, GUI access, management and authentication infrastructure. It is connected to 10 TB of storage for user data and 80 TB of fast NVMe-based storage for the supporting infrastructure.

Local GitLab instance

ARCC hosts a full featured community distribution of the GitLab DevOps platform, permitting UW researchers to collaborate on private code repositories in sensitive projects. This service is provided via a Virtual Machine that is sufficient to support hundreds of active projects. This service is sponsored by the UW EPSCoR.