

Facilities, equipment, and other resources

Expertise

Research Computing at CU Boulder consists of a small group of computational scientists, high-performance computing specialists, and system administrators with the mission to provide leadership in developing, deploying, and operating an integrated cyberinfrastructure. This cyberinfrastructure consists of high-performance computing, on-premise and commercial cloud computing, storage, and high speed networking that supports research, collaboration and discovery. Research Computing contributes to the educational mission of the university by providing training workshops and consultation services for cyberinfrastructure related topics.

Compute

- Research Computing operates the joint RMACC Summit supercomputer, funded by NSF under Grant No. AC- 1532236. The system has peak performance of over 400 TFLOPS. The 472 general compute nodes each have 24 cores aboard Intel Haswell CPUs, 128 GB of RAM and a local SSD. Additionally, the system has 11 GPU nodes with two NVIDIA K80 GPUs each, 5 high-memory nodes with ~2TiB of main memory, and 20 Xeon Phi (“Knight’s Landing”) nodes each with 68 real cores supporting 272 threads. All nodes are connected through a high-performance network based on Intel Omni-Path with a bandwidth of 100 Gb/s and a latency of 0.4 microseconds. A 1.2 PB high-performance IBM GPFS file system is provided. This system is available to CU-Boulder and Colorado State University researchers and collaborators, and 10% of cycles are provided to members of the Rocky Mountain Advanced Computing Consortium.
- Beginning in 2022, Research Computing introduced their third-generation Alpine supercomputer, funded by contributions from the University of Colorado Boulder, Colorado State University, and the University of Colorado Anschutz. Through the second phase of the University of Colorado Boulder contribution to Alpine includes 192 general compute nodes, each with 64

cores aboard AMD Milan CPUs, 256 GB of RAM, and a local SSD. CU's contribution also includes 8 GPU nodes each with 3x NVIDIA A100 GPUs, 8 GPU nodes each with 3x AMD MI100 GPUs, and 12 nodes with 48 cores and 1TB of RAM. Additionally, Alpine is host to 28 48-core nodes and 49 32-core nodes contributed by Colorado State University, as well as 14 general compute nodes, 2 nodes with 1TB of RAM, and 4 3xNVIDIA A100 GPU nodes from Anschutz Medical Campus at the University of Colorado Denver. All nodes are connected through a high-performance network based on Mellanox Infiniband with a bandwidth of 200 Gb/s. A 2 PB high-performance IBM GPFS file system is provided. Additional hardware phases will be added to Alpine each year.

- Research Computing offers a condo computing service that enables researchers to purchase and own compute nodes that are operated as part of a cluster, named "Blanca." The aggregate cluster is made available to all condo partners while maintaining priority for the owner of each node.
- Research Computing provides a 3d-accelerated virtual desktop environment for real-time visualization and rendering using EnginFrame and DCV. This environment is powered by two visualization nodes, each equipped with 2x AMD EPYC 7402 24-core processors, 256GiB memory, and 2x Nvidia Quadro RTX 8000 GPU accelerators. Each accelerator is itself equipped with 48 GiB of high-speed GDDR6 memory.
- Research Computing hosts a free-to-use on-premise cloud service, called CUmulus, which supports cases not well-suited for HPC such as web servers, databases, and long-running services. The CUmulus service includes access to a Virtual Private Cloud (VPC) which provides users with a logically isolated section of the cloud with a small number of outside routable floating IP addresses. Within this VPC customers are given an allocation of CPU cores, memory and storage, which can be used to host virtual machines and volumes to host workloads. The CU Research Computing group would like to acknowledge support of this project from the National Science Foundation (OAC-1925766). In addition to the CUmulus service, Research Computing

can also facilitate researcher access to offsite commercial cloud resources at negotiated rates.

Networking

The current CU Boulder network is a 40 Gbps fiber core with Cat 5 or higher wiring throughout campus. Research Computing has created an 80 Gbps Science-DMZ to connect the RMACC Summit supercomputer to storage and to bring individual dedicated 10 Gbps circuits to various locations as needed. CU Boulder participates in I2 (the Internet 2 higher education, government, and vendor research computing consortium) and is an active member of the Front-Range gigapop and other networks. Research Computing has begun to provide campus researchers with a leading-edge network that meets their needs and facilitates collaboration, high performance data exchange, access to co-location facilities, remote mounts to storage, and real-time communications.

File Transfer

For moving large volumes of data Research Computing has several nodes dedicated to GridFTP file transfer.

The CU Office of Information Technology also offers a file transfer service with a web interface, which provides an ideal way to transfer large files to collaborators. Files are uploaded to a server and a link to download the file is emailed to an on- or off-campus user.

Storage

Each Research Computing user has a 2 GB home directory and a 250 GB projects directory, each of which are backed up regularly. Each RMACC Summit and Alpine user has a 10 TB scratch directory.

PetaLibrary Storage Services

The PetaLibrary is a CU Research Computing service supporting the storage, archival, and sharing of research data. It is available at a subsidized cost to any researcher affiliated with the University of Colorado Boulder. The two main categories of service offered to customers of the PetaLibrary are Active storage for data requiring frequent access, and Archive storage for data that is accessed infrequently. Active data is stored on spinning disk and is accessible to researchers on compute resources managed by Research Computing. Archive storage consists of a two-level hierarchical storage management (HSM) solution, with disk storage for data that is more likely to be accessed and tape storage for data that is less likely to be accessed frequently. The cost for CU researchers is \$45/TB/year for Active and \$20/TB/year for Archive.

Through a collaboration with the CU Libraries, the PetaLibrary can also host the publication and long-term archival of large datasets. The datasets are assigned unique digital object identifiers (DOIs) that are searchable and accessible via the “CU Scholar” institutional repository.

JupyterHub

JupyterHub is a multi-user server for Jupyter notebooks. It provides a web service enabling users to create and share documents that contain live code, equations, visualizations and explanatory text. The CU Research Computing JupyterHub deploys into the RMACC Summit supercomputer and Blanca condo cluster and includes support for parallel computation.

EnginFrame

NICE EnginFrame provides a 3D-accelerated remote desktop environment on Nvidia GPU-equipped compute nodes. Coupled with the proprietary Desktop Cloud Visualization (DCV) VNC server, the Research Computing EnginFrame supports the use of common visualization applications in a typical desktop environment using only a modern web browser.

OnDemand

Open OnDemand is a browser based, integrated, single access portal for CURC high performance computing (HPC) resources. It provides a graphical interface to view, edit, download, and upload files, manage and create job templates for CURC's computing clusters, and access CURC interactive applications (visualization nodes, Matlab, and Jupyter Notebooks), all via a web browser.

Center for Research Data and Digital Scholarship (CRDDS)

The Center for Research Data & Digital Scholarship (CRDDS) is a collaboration between Research Computing and University Libraries, offering a full range of data services for both university and community members. The aim of CRDDS is to provide support to community members on areas related to data intensive research. CRDDS fulfills this mission by providing education and support on such issues as data discovery, reuse, access, publication, storage, visualization, curation, cleaning, and preservation, as well as digital scholarship. CRDDS is located in Norlin Library on Main Campus at CU Boulder.

CRDDS offers many opportunities to students working with data. The expert staff work hand-in-hand with researchers via weekly office hours, one-on-one consultations, and group trainings in programming, data visualization and more. CRDDS serves as a resource for data management, manipulation and publication for trainees working through undergraduate and graduate coursework.

Examples of workshops/trainings CRDDS has offered include:

- High performance computing
- Programming in R
- Programming in Python
- Containerization
- Data mining

