

USC Center for Advanced Research Computing

Description of Facilities and Resources (October 2021)

As the central advanced research computing program at the University of Southern California (USC), the Center for Advanced Research Computing (CARC) provides the computing services and support necessary to enable cutting-edge scientific research. As part of the Information Technology Services (ITS) department, CARC resources are housed within the ITS data center and monitored 24/7. Basic resources are available at no cost to USC researchers. Information about the center and its researchers is available at carc.usc.edu.

HPC Resources

CARC resources include a large-scale HPC cluster system with around 1,300 compute nodes (totaling around 26,000 cores) connected via a 56 gigabits-per-second (Gbps) InfiniBand FDR low-latency interconnection network, and more than 10 petabytes of combined disk storage in four file systems. These 1,300 compute nodes make up the Discovery general-use cluster and the Endeavour condo cluster. Among the nodes, more than 100 of them are equipped with graphics processing units (GPUs) to support accelerated computational simulations as well as advanced data science analysis.

The following table summarizes the HPC resources offered by the CARC:

Category	Function	Note
Discovery general-use cluster	Login nodes	2 40 Gbps nodes
	Data transfer nodes	2 100 Gbps nodes running GlobusConnect
	Compute nodes	~500 nodes, totaling ~11,000 cores
	GPUs	~100 GPUs (K40, V100, P100)
	Large memory nodes	3 nodes with 1 TB of memory
Endeavour condo cluster	Login nodes	2 40 Gbps nodes
	Compute nodes	~800 nodes, totaling ~15,000 cores
Network	Interconnection	InfiniBand FDR (56 Gbps)
Storage file systems	/home1	100 GB/user NFS/ZFS file system
	/project	10 TB/project(s) ZFS/BeeGFS parallel file system
	/scratch	10 TB/user ZFS/BeeGFS parallel file system
	/scratch2	10 TB/user ZFS/BeeGFS parallel file system

Network

To support advanced computing at the university, USC recently upgraded the network capacity of individual research buildings to 100 Gbps. The CARC has also deployed multiple 100 Gbps links between the ITS data center and the rest of the campus, creating an aggregate bandwidth of up to 400 Gbps.

CARC Accounts

CARC accounts are available at no cost to USC faculty, research staff, and graduate students with the approval of their faculty advisor. USC researchers, undergraduate students, and external collaborators with iVIP accounts may be added as members to CARC accounts. The CARC also offers HPC Secure Data Accounts (HSDA) for researchers working with legally restricted data. Specialized resources for HSDA accounts provide a secure environment to meet HIPAA requirements. Data stored in these accounts will be encrypted automatically upon upload and at rest.

Condo Cluster Program (CCP)

The Condo Cluster Program (CCP) is a service available to USC researchers that require dedicated resources for their work. The CCP gives researchers the convenience of having their own dedicated compute nodes, without the responsibility of purchasing and maintaining the nodes themselves. The CCP operates on two different models – an annual subscription model and a traditional system purchase model – to provide researchers with flexible and efficient options for their resources.

Compute nodes leased through the CCP form the Endeavour condo cluster and are configured for exclusive use by the condo owner. System policies for condo partitions are based on each condo owner's specialized requirements and can include optimization features not available under general use partitions. All administrative work for condo nodes is covered as part of CARC's basic services.

Research Computing Software

Researchers can install software packages or develop their own code within their project's allocated storage. The CARC provides an array of scientific software packages, both licensed and open-source, for use on the cluster, such as multiple compiler options supporting OpenMP and MPI programs as well as applications like Python, R, Anaconda, and MATLAB. These resources are available to USC researchers at no cost.

Support and Training

User support for advanced computing at USC is primarily available through the CARC, which consists of a systems team, cloud/DevOps team, and research support team. The staff has expertise in storage system administration, software development, science and engineering applications, high-performance and parallel programming, and cloud computing solution development. The CARC also works with researchers who do not require HPC-level performance but need help developing or optimizing their code, improving their computing productivity by learning better programming languages, scripting, and so forth.

The CARC offers regular onboarding sessions, office hours, introductory and advanced training workshops, and in-person and remote consultations. It also hosts vendor and national supercomputing center training at USC. As XSEDE Campus Champions, CARC staff can connect researchers with high-end computing needs to national supercomputing centers to scale their workflows and reach other researchers nationwide.