

Facilities Managed by Research Cyberinfrastructure Center (RCIC) at UC Irvine

Last Updated: June 2021

The Research Cyberinfrastructure Center (RCIC) at UC Irvine has seven permanent staff (several Ph.Ds and M.S.) to develop, deploy, and maintain data and computing infrastructure for academic research. The key infrastructure includes:

1. High-performance/High-Throughput Computing Clusters (HPC² and HPC³).
2. A rich set of domain application software installed and maintained with more than **260** distinct applications available to UCI researchers.
3. Single-copy, high-performance, parallel file systems based on BeeGFS. Directly connected via Infiniband networking to the clusters, these provide an aggregate of Four Petabyte (4PB) across six active file systems (DFS2, DFS3a, DFS3B, DFS4, DFS5, DFS6).
4. Dual-copy, 2PB, scalable file system based on IBM Spectrum Scale for research data. Using standard protocols, data can be accessed securely from any device on the UC Network (Campus Research Storage Pool, CRSP).

All resources are available to the UCI research community. Network connectivity is through the standard UCI Network. A small number of systems are connected to LightPath (UCI's science DMZ network). Both networks are managed by the UCI Office of Information Technology (OIT).

Clusters managed by RCIC

The HPC2 and HPC3 production clusters are managed by RCIC. All parallel file systems (DFSx, CRSP) are available on both. The HPC2 cluster is built from still-usable hardware of the former HPC cluster, which was decommissioned in January of 2021.

In the second quarter 2020, HPC3, UCI's flagship cluster, entered early operation. Full production was achieved in Fall 2020. This resource is available to all researchers at UCI. It can be extended via hardware purchase or core-hours purchase. Faculty, upon request, are granted about 200,000 core hours per calendar year. These hours are made possible by UCI-purchased hardware. HPC3 is expected to grow over time with a mixture of CPU-types. It is multi-vendor with pure compute and compute + GPU nodes. All systems are interconnected with 10Gbps Ethernet and 100Gbps EDR Infiniband. All nodes are configured with at least 4GB/core of memory. Configuration as of June 2021 is as follows:

- 149 Nodes Total, 6952 cores
- 13 GPU nodes – each with 4 Nvidia V100 GPUs (total of 52 GPUs)
- Most nodes have at least 40 cores

Most nodes have a standard memory of 4.0GB/core (standard node), but there are other node configurations including 9GB/core (large memory), 18GB/core (huge memory), and 36GB/core (maximum memory). A collection of application software is deployed on HPC3 using a rigorous and standardized software build and deployment process.

The HPC2 facility is a transitional cluster and is available only to researchers at UCI who owned existing HPC hardware that still had usable life but was not compatible with HPC3. There are no plans to expand HPC2. This cluster is interconnected with 40Gigabit Infiniband (QDR) and 1 Gigabit Ethernet. Cores are of mixed vendor, age, frequency, and capability. The software stack is identical to the one on HPC3.

Application Software Stack

RCIC manages, installs, and updates a rich collection of scientific software spanning a very large number of disciplines including, biology, chemistry, physics, medicine, engineering, social sciences, humanities, and data sciences. More than 260 distinct software modules (which are not available through standard OS channels) are handled on HPC3. Environment modules enable users to load different versions when needed. When taking software with multiple versions and with multiple additional packages (for Perl, R and Python) more than 1700 distinct software packages are built and maintained by RCIC.

Single-Copy Parallel File Systems

Large-scale data is handled by RCIC across six distinct parallel file systems that are named DFS2, DFS3a, DFS3b, DFS4, DFS5, and DFS6. All systems utilize BeeGFS layered on ZFS for scale and performance. BeeGFS is regularly updated with one to two major updates per year. Storage is funded by researchers through a standard recharge system where a user purchases capacity for a 5-year period of time (currently \$100/TB/5-years). The DFS systems are directly available only to the clusters and a handful of data transfer nodes.

Dual-Copy Campus Research Storage (CRSP)

UCI has recognized the importance of high-quality storage for research data. The Campus Research Storage Pool (CRSP) is 1 Petabyte of usable capacity IBM Spectrum-Scale system where researchers can access their research data through the UCI Network. Performance is quite good so that data may be used in-place. Data is immediately replicated into two distinct data centers on campus. Significant redundancy in hardware results in very high data availability. Since entering production in early 2019, CRSP has not had a full outage.

Costs for Adding Storage, Processing, GPU Capability

Resources can be expanded via grant (or other funding) purchases. For computing, grants may either purchase equipment (condo) and pay a one-time integration charge, or purchase a prepaid time allocation (cloud-like). RCIC does not post-bill for resources.

RCIC and its advisory committee revisit “standard” hardware configurations annually. Only standard configurations are supported. For HPC3, when a condo node is added, the number of core hours than node could deliver in a year is added to the owner’s account. The owner (and their students/colleagues) can access all nodes of similar configuration. Accounting is used to ensure that the number of cycles (core hours or GPU-hours) available to a condo owner over a fixed period of time is *equivalent to or better than dedicated access to specific hardware*.

The following figures can be used for grant-planning purposes, actual costs of hardware are commodity and require formal quotations at the time of purchase. For grant applications, the most-recent quotations can be supplied. Due to the pricing volatility of commodity computing, quotations at grant submission time will almost certainly expire prior to grant award. Standard nodes may have parts updated (e.g. CPU change or GPU change) as time evolves.

Standard Node Configurations for new purchase (Valid for Calendar Year 2021)

Description	Brief Specification	Approximate Cost (pre-tax)
CPU Node	2x Intel 6248 (24 Core, 2.4Ghz. 48 cores total), 1 x Infiniband IB (EDR), 2 X 10/25GigE SFP, 1 x 480GB SSD, 1x1.92TB SSD (scratch), 192 GB Memory. 3 Yr. Warranty	\$8200
GPU Node	Same as CPU Node + 4 x Nvidia A100 GPU with 32GB of High-bandwidth Memory	\$55000
384GB	Expand Memory from 192GB to 384GB	\$2500
768GB	Expand Memory from 192GB to 768GB	\$6500
CPU Warranty	Extend Warranty on CPU node from 3 years to 5 years	\$500
GPU Warranty	Extend Warranty on GPU node from 3 years to 4 years	\$1500

Table 1 - Approximate cost of CPU/GPU nodes for equipment purchase. These are standard configurations.

The following table summarizes costs and time-commitments for all services.

Item	Brief Description	Time Commitment	Cost
Condo Node Racking	Racking a condo node into HPC3. RCIC purchases switch ports, cables from this fee	Once (per node)	\$1000
CPU Core-Hour	Cost to purchase a core-hour of computing from RCIC-purchased hardware	Expires after 1 Year. Minimum Purchase: 10000 core hours	\$0.01/ Core-hour
GPU Hour	Cost to purchase a GPU hour from RCIC-purchased hardware	Expires after 1 Year. Minimum Purchase: 500 GPU Hours	\$0.32/GPU-Hour
1 TB DFS	1 Terabyte of quota on a DFS system. Single copy of data. Accessible only from the clusters	5 years	\$100/TB/5 years
1 TB CRSP	1 Terabyte of quota on CRSP. Dual copies of data. Accessible from Devices on UCI net and the clusters	1 year	\$60/TB/year

Table 2 - Costs of Services Supported by RCIC