



---

# RESEARCH COMPUTING & DATA SERVICES AT USC

CENTER FOR ADVANCED RESEARCH COMPUTING (CARC)

VITERBI SCHOOL OF ENGINEERING  
NEW FACULTY ORIENTATION AUGUST 17, 2023





---

# ABOUT CARC

<https://www.carc.usc.edu/>

The computational expertise in high-performance computing of the USC Center for Advanced Research Computing (CARC) has been a vital resource in USC research community and contributes to improved research productivity and superior outcomes, driving USC's research excellence forward.

**CARC supports USC's mission by providing advanced research cyberinfrastructure and the computational expertise necessary to enable cutting-edge scientific research.**

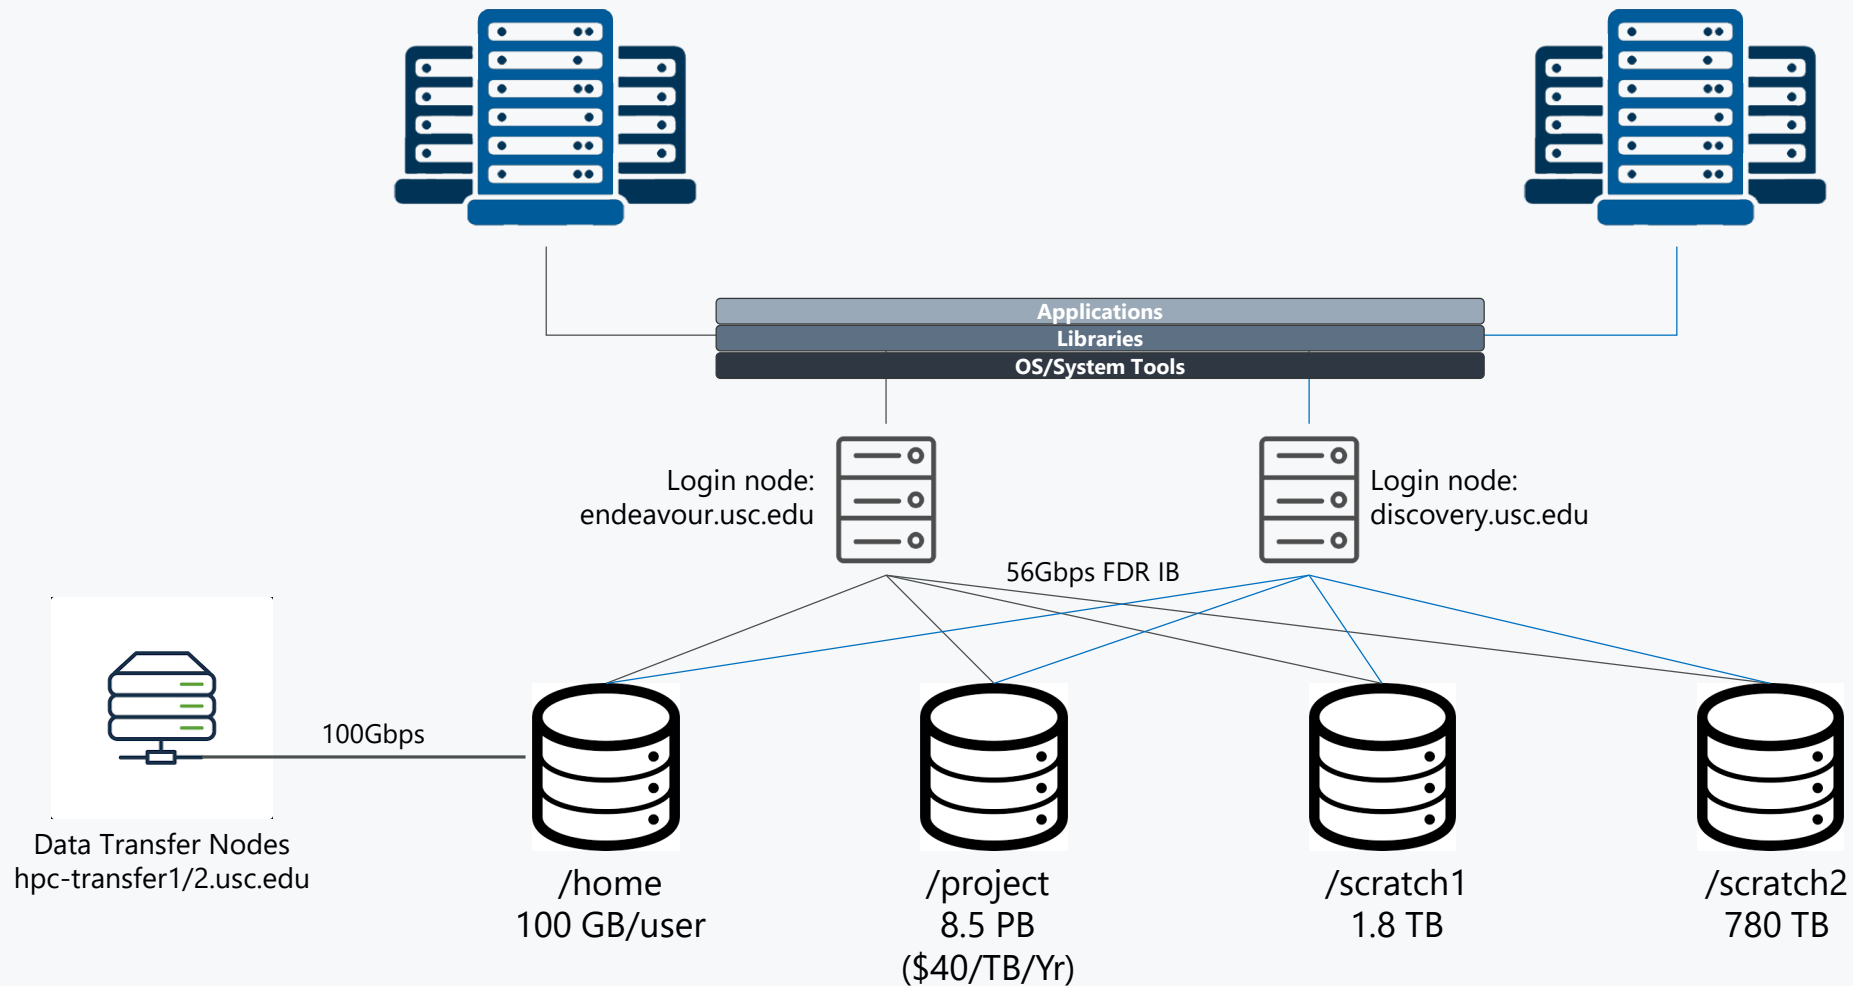


# CARC SYSTEMS OVERVIEW

CARC systems include the Endeavour condo cluster as well as the Discovery shared cluster

Endeavour (Condo Cluster)

Discovery (Shared Cluster)



# HPC RESOURCES

The following table summarizes the resources at CARC.

Category	Function	Descriptions
<b>Discovery general-use cluster</b>	Login nodes	2 x 40 Gbps nodes
	Data transfer nodes	2 x 100 Gbps nodes running GlobusConnect
	Compute nodes	~600 nodes, totaling ~22,000 cores
	GPUs	~360 GPUs (A100, A40, K40, V100, P100)
	Large memory nodes	4 nodes with 1 TB of memory
<b>Endeavour condo cluster</b>	Login nodes	2 x 40 Gbps nodes
	Compute nodes	~900 nodes, totaling ~22,000 cores
	GPUs	~180 GPUs (V100, A40, A100)
<b>Network</b>	Interconnection	InfiniBand FDR (56 Gbps) – Soon to be upgraded to NDR
	Science DMZ	DTN w/ perfSONAR
<b>Storage file systems</b>	/home1	280TB total, 100 GB/user ZFS/NFS parallel file system
	/project	9PB ZFS/BeeGFS parallel file system
	/scratch1&2	2.6PB total ZFS/BeeGFS parallel file system
<b>Artemis cloud platform</b>	Compute nodes	14 nodes, totaling 896 cores
	GPUs	6 x A40 GPUs

Recognizing the need for a significant upgrade in USC's current HPC infrastructure to meet the demands of future research, CARC is actively engaged in collaborative efforts with multiple departments across the university, focusing on planning and developing next-generation HPC systems that will deliver the necessary computing power to propel USC's research endeavors forward. These concerted efforts aim to provide optimal support for the the Frontiers of Computing initiative, ensuring that USC remains at the forefront of cutting-edge research and innovation.

# CARC SYSTEM USAGE STATUS

Discovery HPC cluster usage status in a quick glance

# OF PROJECTS

900+

## On Discovery

- 1000+ PI's
- 5,000 users
- 300M SU's have been allocated
- 160M SU's used on the system since inception

# OF NODES ON DISCOVERY

600

## System Specs

- Total 22K cores
- 300+ applications
- 360 GPUs
- 148 old compute nodes will be decommissioned in FY24

FY23 COMPUTE PURCHASE

\$0

## New System Purchase: ~~None~~ HPC Network Upgrade

- Current: FDR (56Gbps)
- Upgrading to:
  - NDR (200Gbps)



---

# CCP: CONDO CLUSTER PROGRAM

The Center for Advanced Research Computing (CARC) launched the Condo Cluster Program (CCP) in December 2020 to allow researchers a flexible way to purchase computing resources for their own dedicated use.

**The CCP has two pricing models:**

## **Annual Subscription Model**

- Allows research groups to subscribe to their selected number of compute and storage resources on a yearly basis
- Compute resources can be requested via CARC User Portal
- Allocated nodes get provisioned automatically within a week

## **Traditional 5-year System Purchase Model**

- A useful option when research groups need to make a bulk purchase using a research grant or departmental budget
- Compute/GPU system configurations by CARC
- System purchases can be requested via CARC User Portal



# STATE OF THE NEW CONDO CLUSTER PROGRAM

Current usage for CARC's Endeavour condo cluster

# OF PROJECTS:

## 61

### Contributions:

- Viterbi: 15 PI's, 393 nodes
- Dornsife: 26PI's, 256 nodes
- Keck: 13 PI's, 131 nodes
- Others: 2 PI's, 3 nodes

- Not including subscription nodes

# OF NODES IN USE:

## 900

### System utilization:

- 22K cores
- 180 GPUs
- 42 subscription nodes
- 257 old compute nodes will be decommissioned in FY24

FY21-22 PURCHASE:

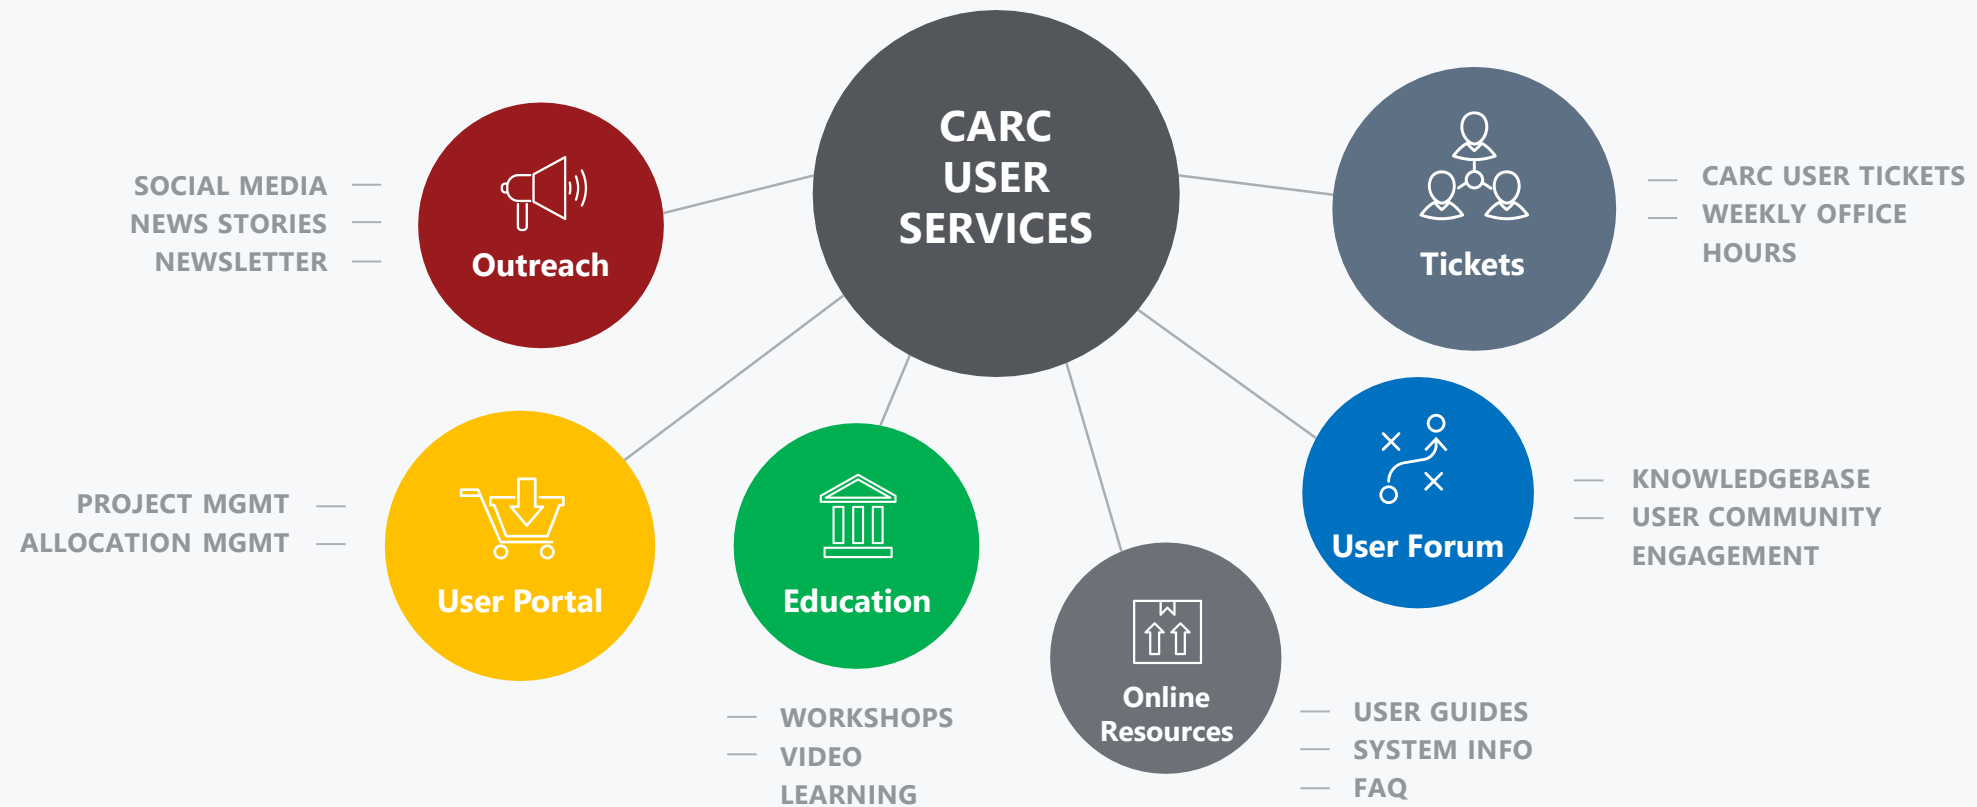
## \$660K

### New nodes purchased:

- 18 CPU nodes
- 9 GPU nodes (12 x A40 / 12 x A100)
- DOE cluster (18-node)

# ADVANCED RESEARCH COMPUTING **USER SERVICES**

The Center for Advanced Research Computing (CARC) offers comprehensive user support services





---

# CARC HIGHLIGHTS

The resources, services, and achievements that set us apart.



## Research Collaborations:

- NSF Campus Compute (\$400K) - Hybrid Computing System Development
- NSF Regional Network (\$1M) - Science DMZ R&E Network Deployment
- NSF Regional Computing (\$1M) – Leading Research Computing Alliance in SoCal
- Cryo-EM project for USC & Amgen –Development of Computational Ecosystem



## Education & Outreach:

- More than 20 workshop classes and summer bootcamp
- ITP 450: High-Performance Computing for Applied Machine Learning
- NSF CyberTraining proposal (\$300K) development with AME faculty



## Advanced Cyberinfrastructure:

- Discovery shared HPC cluster system & Endeavour Condo cluster program
- Artemis: virtual computing platforms and cloud solution (NSF Funded)
- High-performance HPC network upgrade to 200Gbps
- 10+PB data storage capacity



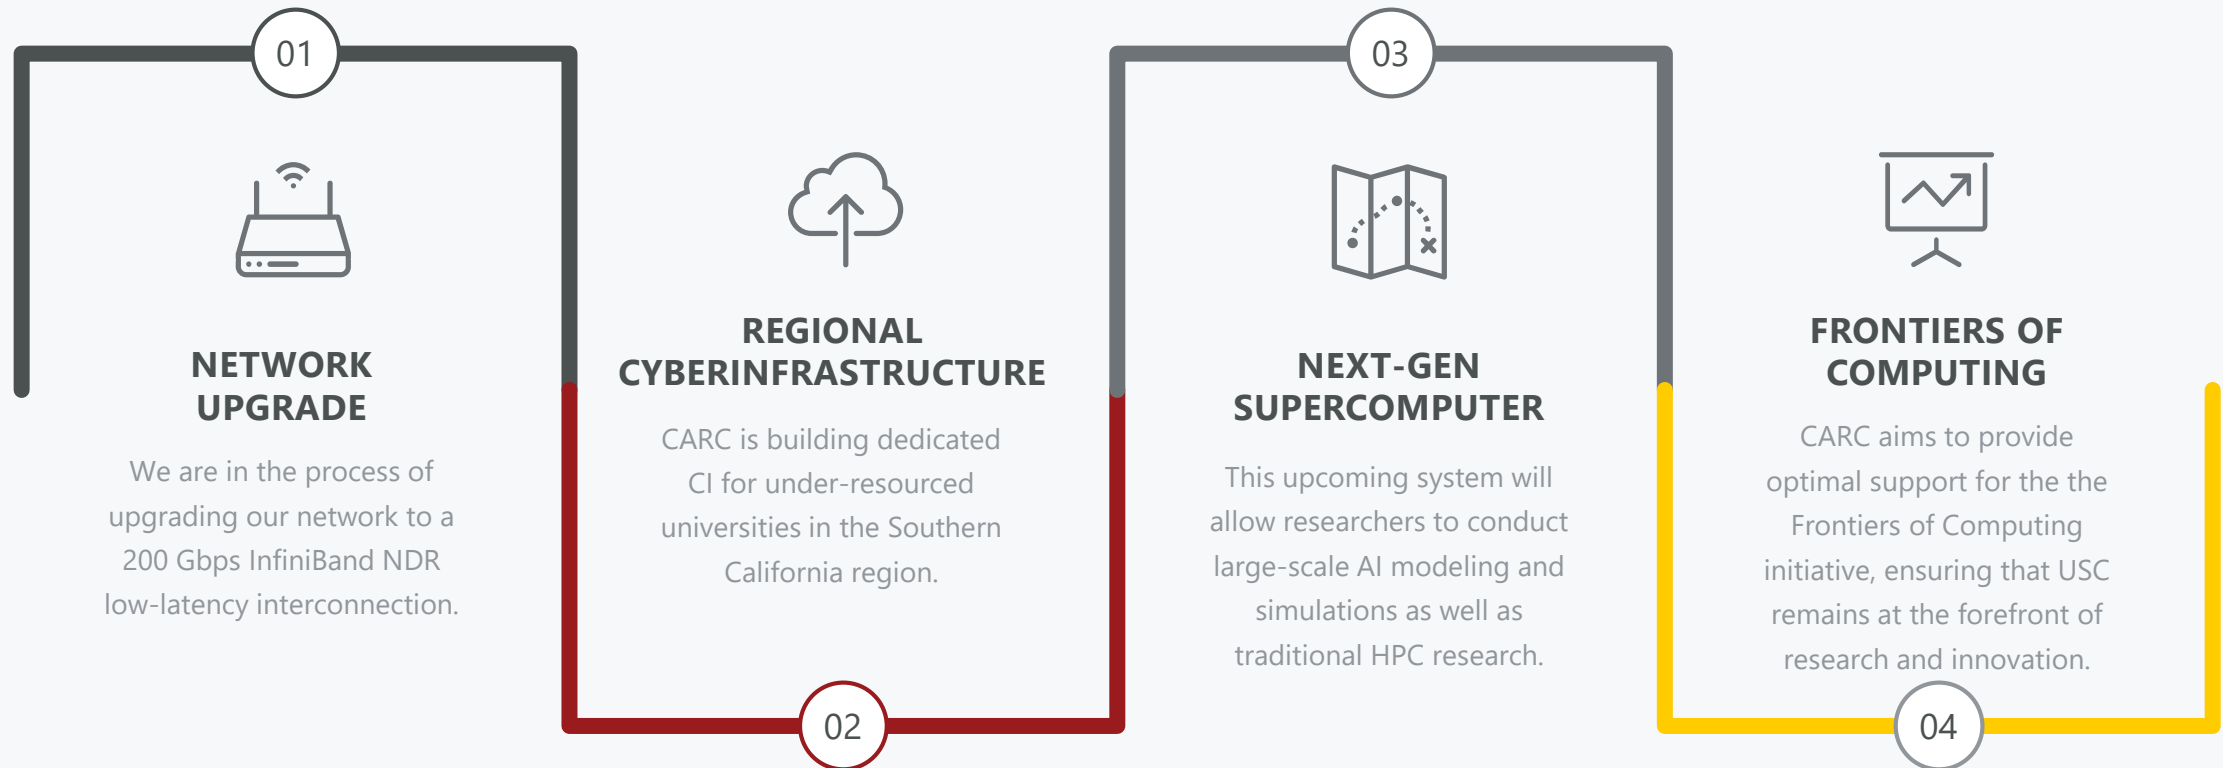
## Industry Partnership:

- Samsung Semiconductor, Inc. – Full NVMe storage solution (2+PB)
- VAST Data – Advanced FS testbed
- Nvidia – Early access to Grace-Hopper next-gen system design & NSF proposal



# USC RESEARCH COMPUTING & DATA ROAD MAP

The future of CARC looks bright with plans for future improvements, outreach, and collaborations.



---

# THANK YOU & FIGHT ON!



**USC** University of  
Southern California

