

HPC Capabilities at Research Intensive Universities

Purushotham (Puri) V. Bangalore

Department of Computer and Information Sciences and

UAB IT Research Computing

UAB THE UNIVERSITY OF
ALABAMA AT BIRMINGHAM

Knowledge that will change your world

UAB HPC Resources

- 24 nodes (192 cores) with Intel Quad-Core E5450 3.0 GHz CPUs and 16 GB per node connected with DDR InfiniBand
- 48 nodes (576 cores) Intel Westmere 2.66 GHz CPUs with 24 GB per node connected with QDR InfiniBand (NIH SIG grant led by Dr. Allison)
- All nodes connected to 240TB of Lustre parallel file system using InfiniBand network on DDN S2A9900
- 20 TB of home directory space on Hitachi SAN
- 12 Dell 720xd systems each with 16 cores, 96GB RAM, and 36TB of storage, creating a 192 core, 432TB virtual compute and storage (cloud) fabric

Alabama Supercomputing Authority (ASA) Resources

- SGI UV
 - Compute Nodes – 256 Xeon E5-4640 CPU cores with 4 TB of shared memory
 - Frontend – 16 Xeon E5-2667 CPU cores with 64 GB RAM
- DMC – 1800 CPU cores + 16 GPUs
 - 40 nodes – 2.3 GHz quad-core AMD Opterons and 64 GB RAM
 - 96 nodes – 2.26 GHz Intel quad-core Nehalem processors
 - 40 nodes – 2.3 GHz AMD 8-core Opterons and 128 GB RAM
- SGI Altix
 - 162 Intel Itanium 2 processors with 1340 GB of shared memory
- Storage
 - Scratch - 19 TB Panasas file system
 - Home directory – 47 TB in a separate Panasas file system

XSEDE HPC Resources

- TACC Stampede system – 10 Pflop/s Cluster based on 6,400+ Dell PowerEdge server nodes, each outfitted with 2 Intel Xeon E5 (Sandy Bridge) processors and an Intel Xeon Phi Coprocessor
- TACC Lonestar Cluster – consists of 1,888 compute nodes, with two 6-Core processors per node, for a total of **22,656 cores**. It is configured with 44 TB of total memory and 276TB of local disk
- NCIS Kraken - Cray XT5 with 9,408 compute nodes, each compute node has two six-core 2.6 GHz AMD Opterons for a total of **112,896 cores**.
- Georgia Tech Keeneland system – consists of 264 HP SL250G8 compute nodes, each with two 8-core Intel Xeon E5 (Sandy Bridge) processors, three NVIDIA M2090 GPU accelerators, for a total of 264 nodes, 528 CPUs and 792 GPUs.

XSEDE HPC Resources (continued)

- SDSC Gordon – flash-based supercomputer designed for data intensive applications that has 1024 compute nodes and 64 I/O nodes. Each compute node contains two 8-core 2.6 GHz Intel Xeon E5 (Sandy Bridge) processors and 64 GB RAM, each I/O node contains two 6-core 2.67 GHz Intel X5650 (Westmere) processors, 48 GB of RAM, and 16 300GB solid state drives.
- SDSC Trestles – consisting of 324 compute nodes, Each compute node contains four sockets, each with a 8-core 2.4 GHz AMD Magny-Cours processor, for a total of 32 cores per node and **10,368 total cores** for the system
- PSC Blacklight – 37 Tflop/s SGI UV 1000 cc-NUMA shared-memory system with 256 blades (**4096 cores**), each blade has 2 Intel Xeon X7560 (Nehalem) eight-core processors

XSEDE HPC Resources (continued)

- IU Mason – a large memory computer cluster configured to support data-intensive, high-performance computing tasks. Mason consists of 16 HP DL580 servers, each containing four Intel Xeon L7555 8-core processors and 512 GB of RAM.
- Mason is intended for use by researchers using genome assembly software (particularly software suitable for assembly of data from next-generation sequencers), large-scale phylogenetic software, or other genome analysis applications requiring large amounts of computer memory.
- National Center for Genome Analysis Support (NCGAS.org) – provides access to the above XSEDE resources for NSF-funded researchers at US-based accredited institution of higher education or federal research lab.

Other HPC Resources

- Open Science Grid (OSG) – suitable for high throughput workflows
- Application-specific Gateways
 - Biodrugscore: A portal for customized scoring and ranking of molecules docked to the human proteome
 - Computational Chemistry Grid (GridChem)
 - ParamChem Gateway
 - Network for Computational Nanotechnology and nanoHUB
 - National Biomedical Computation Resource
 - The iPlant Collaborative Agave API
 - Integrated database and search engine for systems biology (IntegromeDB)
 - ROBETTA: Automated Prediction of Protein Structure and Interactions
 - Providing a Neuroscience Gateway
 - CIPRES Portal for inference of large phylogenetic trees
 - Many more.... Full listing at: <https://www.xsede.org/gateways-listing>

How to access XSEDE HPC resources?

- Go to <http://www.xsede.org>
- Learn about XSEDE resources and services
- Sign up for an XSEDE User Portal (XUP) account
- Get an allocation¹ (startup, education, or research)
- Access your allocated resources
- Transfer your files
- Manage your Computing Environment
- Sign up for training
- Get Help
- Access and manage XSEDE resources with Mobile User Portal


¹To apply, the PI must be a researcher or educator at a U.S.-based institution, including federal research labs or commercial organizations. investigators with support from any funding source, not just NSF, are encouraged to apply.

Academic institutions in the Top 500 supercomputers list

| Rank | Site | Cores | Rmax (TFlop/s) | Rpeak (TFlop/s) |
|------|-------------------------------------|---------|----------------|-----------------|
| 6 | TACC/Univ. of Texas | 462,462 | 5,168.1 | 8,520.1 |
| 28 | Purdue University | 77,520 | 961.9 | 1,341.1 |
| 30 | NICS/Univ. of Tennessee | 112,800 | 919.1 | 1,173.0 |
| 46 | Indiana University | 31,288 | 597.4 | 1,000.6 |
| 53 | University of Southern California | 10,920 | 531.6 | 690.4 |
| 76 | Rensselaer Polytechnic Institute | 32,768 | 357.8 | 419.4 |
| 102 | UCSD/San Diego Supercomputer Center | 16,160 | 285.8 | 336.1 |
| 115 | Clemson University | 5,376 | 262.4 | 339.9 |
| 119 | TACC/Univ. of Texas | 22,656 | 251.8 | 301.8 |
| 171 | University of Rochester | 16,384 | 189.0 | 209.7 |
| 176 | Purdue University | 10,368 | 186.9 | 215.7 |
| 240 | University of Colorado | 15,648 | 152.2 | 175.3 |
| 243 | University of Southern California | 20,925 | 149.9 | 196.5 |
| 302 | University of Chicago | 6,880 | 130.2 | 143.1 |
| 331 | University of Chicago | 17,856 | 125.8 | 150.0 |
| 351 | Virginia Tech | 8,320 | 120.4 | 238.2 |
| 393 | Louisiana State University | 7,040 | 110.9 | 146.4 |
| 397 | NICS/Univ. of Tennessee | 9,216 | 110.5 | 157.5 |
| 399 | University of Miami | 7,520 | 110.0 | 146.9 |
| 497 | Clemson University | 13,616 | 96.9 | 123.6 |

Top 500 Supercomputers Performance

Performance Development



UAB HPC Resource Requirements

- Computing
 - Few thousand CPU cores
 - Some nodes with larger memory
 - Some nodes with accelerators (GPU or Phi)
 - Some nodes with additional storage (SSD and/or disks)
- Storage
 - Home directory
 - Shared project
 - Scratch
 - Archival
- Networking
 - FDR InfiniBand and 10 Gigabit Ethernet within the cluster
 - 10 Gigabit Ethernet connection to access the cluster

Sample Configuration

- 108 nodes with
 - Two eight-core CPUs at 2.6 GHz
 - 96 GB RAM
 - FDR InfiniBand and 10 Gigabit Ethernet
- 10 nodes have an accelerator (GPU or Phi)
- 10 nodes have extra storage (around 500 TB raw)
- Total core count = 2,048
- 500 TB of home directory storage
- Approximate cost around \$1M

Funding Opportunities

- NSF Major Research Instrumentation
 - NSF-supported fields of science, math, and engineering
 - Two proposals per organization for acquisition
 - 175 awards for budget requests in the range \$100,000–\$4M
 - 30% cost-sharing required
 - Deadline: January 23, 2014
- NIH ORIP Shared Instrumentation Grant (S10)
 - Budget range: \$100,000–\$600,000
 - Deadline: March 2014
- NIH ORIP High End Instrumentation Grant (S10)
 - Budget range: \$750,000–\$2,000,000
 - Deadline: September 2014

Interested?

Contact: puri@uab.edu