7-1-2020

Rosen Center for Advanced Computing: Facilities, Equipment, and Resources

X Carol Song
Purdue University, carolxsong@purdue.edu

Purdue University Office of Research and Partnerships

Research Computing and Information Technology at Purdue
(Last updated: June 18, 2020)

Please contact Dr. X. Carol Song for more information regarding support for your proposal:
Dr. Carol Song, Director of Scientific Solutions, RCAC, 496-7467, carolxsong "at" purdue.edu

The Rosen Center for Advanced Computing (RCAC) (http://www.rcac.purdue.edu) is the research computing arm of Information Technology at Purdue (ITaP), the University's central IT organization. RCAC provides a large and diverse set of high-performance computing, data-intensive and cloud computing resources, high-speed network connections to national research wide-area networks, and large data storage and archival systems to the broad research communities at Purdue and elsewhere. RCAC also provides resources, expertise, services, and software development efforts to the national research community through partnerships with national cyberinfrastructure projects such as the XSEDE (formerly TeraGrid), Open Science Grid (OSG), and the Science Gateways Community Institute (SGCI).

RCAC offers a broad range of technical support services to Purdue researchers, including user support, consultation, training, and documentation; software installation; computation optimization and scaling; large-scale data management; capacity planning; and cluster deployment. RCAC researchers and software engineers also collaborate with faculty and students to create solutions and tools through software development, integration, and hosting services. The RCAC staff are provided with office space in Young Hall at Purdue University.

The paragraphs below provide additional details of Purdue's centrally supported research computing resources.

Advanced Campus Cyberinfrastructure

RCAC operates and maintains three main data centers across the campus. The main research computing data center, an ~8,000 ft^2 building with 2.5 MW of power and 350 tons of cooling capacity, houses Purdue's centrally managed community clusters, virtual machines and servers, data storage systems, and high-speed network equipment. All of the computing clusters and storage systems are connected via 160 gigabits per second (Gbps) links to wide area research networks such as Internet2 and XSEDEnet.

Computing and Data Resources

RCAC operates several large, state-of-the-art computing clusters as well as data storage and other resources to serve the research and education community at Purdue campuses and nationwide. These systems support scientific computation and data analysis in many research domains ranging from natural sciences, engineering to social sciences. Most of Purdue's high-performance computing (HPC) clusters are accessible via remote desktop tools and Jupyter Hub. The RCAC operated systems include:

High-performance computing

- *Bell*, the latest Community Cluster to join the Purdue campus HPC resources, is optimized for running traditional, tightly-coupled science and engineering applications. Bell was built through a partnership with Dell and AMD during the summer of 2020. Bell consists of Dell compute nodes with two 64-core AMD Epyc 7662 "Rome" processors (128 cores per node) and 256 GB of memory. All nodes have 100 Gbps HDR Infiniband interconnect and a 6-year warranty.
- *Weber* is a high-performance computing cluster specifically designed for data, applications, and research that are covered by export control regulations such as EAR or ITAR or require compliance

with NIST SP 800-171. Built through a partnership with HP and AMD, Weber entered into production in August 2019. Weber consists of HP compute nodes with two 10-core Intel Xeon-E5 "Haswell" processors (20 cores per node) and 64 GB of memory. All nodes have 56 Gbps EDR Infiniband interconnect.

- *Gilbreth* is a GPU cluster built through a partnership with Dell and Nvidia in December 2018. Gilbreth nodes consists of various compute nodes with 128 to 192 GB of system memory. Each node has two GPUs per node with a total of 56 P100 GPUs and 44 V100 GPUs in the system. For training large models, three nodes are available with 768 GB of RAM, 8 TB of local flash storage, and four 32 GB V100 GPU accelerators with Nvlink. Gilbreth nodes share access to a 2.5 PB Lustre filesystem and a 150TB all-flash shared storage system.
- *Brown* is designed and optimized for communities running traditional, tightly-coupled science and engineering applications. Brown was built through a partnership with Dell and Intel in October 2017. Brown consists of Dell compute nodes with two 12-core Intel Xeon Gold "Sky Lake" processors (24 cores per node) and 96 GB of memory. Brown currently has 550 nodes, which totals 13,200 cores on the cluster. All nodes have 100 Gbps EDR Infiniband interconnect and a 5-year warranty.
- *Halstead* is a community cluster built through a partnership with HP and Intel in November 2016. It consists of 508 HP compute nodes, each with two 10-core Intel Xeon E5 processor (20 cores total per node) and 128 GB memory. All nodes are connected through the EDR Infiniband interconnect and to a Lustre parallel file system built on Data Direct Networks' SFA12KX EXAScaler storage platform. This cluster will be in operation through December 2021.
- *Rice* is a community cluster built through a partnership with HP and Intel in April 2015. Rice consists of HP compute nodes, each with two 10-core Intel Xeon-E5 processors (20 cores per node) and 64 GB of memory. Rice has a total of 13,200 processor cores and clocked in with a sustained, measured maximum speed of 407.1 teraflops and a peak performance of 549.1 teraflops in testing. All nodes have Mellanox 56 Gb FDR Infiniband interconnect and a Lustre parallel file system built on Data Direct Networks' SFA12KX EXAScaler storage platform.
- *Snyder* is a big-memory system optimized for data intensive applications requiring large amounts of shared memory per node, such as those used in life sciences research. Snyder consists of more than 70 HP compute nodes, each with two 10-core Intel Xeon-E5 processors (20 cores per node) and various memory configurations of 256 GB, 512 GB, and 1 TB. All nodes have 40 Gbps Ethernet connections.
- *Hammer* is a high-throughput computing system that is optimized for scientific applications running loosely-coupled computation tasks. Initially built through a partnership with HP and Intel in April 2015, Hammer was expanded in late 2016 with 40 new nodes, each with two hyper-threaded 10-core Intel Xeon E4-2660 v3 processors (40 cores per node) and 128 GB of memory. All nodes have 25 Gbps Ethernet connections. The number of nodes in Hammer will be expanded annually.
- *Scholar* is a cluster designed for instructor and student use. This system is open to Purdue instructors from any field whose classes include assignments that could make use of supercomputing or data science. Account management in Scholar is integrated with the university's registration system. The system hardware consists of HP compute nodes with two 8-core Intel Xeon-E5 processors (16 cores per node), 32 GB of memory, and a 500 GB system disk. All nodes have 56 Gbps FDR Infiniband connections. Scholar also provides seven large interactive frontends to teach data science, GPU frontends, and a queue providing access to Nvidia V100 GPUs.
- *Data Workbench* is an interactive computing environment for non-batch big data analysis and simulation. The Data Workbench consists of HP compute nodes with two 8-core Intel Xeon processors (16 cores per node) and 256 GB of memory. All nodes are interconnected with 10 Gigabit Ethernet.

Available tools include Jupyter, R Studio, and Remote Desktop. This resource is part of the Purdue research computing cyberinfrastructure with high-speed network access to the storage systems.

Network connectivity

All Purdue clusters and storage systems are connected with redundant 160 Gbps Ethernet links to the research network core. The campus systems' wide-area connection to the Indiana GigaPOP and major research networks including XSEDEnet, OmniPOP, and ESNET is 100 Gbps.

Public-facing research computing servers and data transfer nodes are accessible via IPv6. All research computing systems support outbound IPv6 in a dual stack mode.

The dedicated campus research network is built in the Science DMZ model featuring a high-speed core, Globus data transfer nodes, a set of PerfSONAR monitoring nodes at key points around the network, and distinct security policies from the campus at large. Currently, PerfSONAR is in production at the core of the research network, but ten PerfSONAR nodes are being deployed at key points of the network infrastructure including the research core, the primary campus routers, and the wide area network entry points.

Data Storage

Purdue operates and maintains a multi-tier storage system to meet the needs of researchers conducting computation, data analysis, data management and sharing, and data archiving.

- *Research Data Depot*: The Data Depot is a high-capacity, fast, reliable, and secure data storage service designed, configured, and operated for the needs of Purdue researchers in any field and shareable with both on-campus and off-campus collaborators. The Data Depot uses an enterprise-class GPFS storage solution with an initial total capacity of over 2 PB and will continue to expand to meet growing demands. This storage is non-purging, redundant, and reliable; features regular snapshots; and is accessible from all Purdue research systems. Built on Data Direct Networks' SFA12k storage platform, the Data Depot has redundant storage arrays in multiple campus datacenters for maximum availability. Multiple data transfer methods are supported, including SCP/SFTP, CIFS/SMB (as Windows drive), and Globus data transfer for fast and easy unattended transmission of large amounts of data between local systems or to/from national resources and labs.
- *Scratch Storage*: Each computing cluster has a default Lustre parallel file system to provide work-area storage optimized for a wide variety of job types and designed to perform well with data-intensive computations while scaling well to large numbers of simultaneous connections. Scratch storage currently consists of 1.4 to 3.4 PB of redundant, high-availability disk space. All scratch tier resources are high-performance, large capacity, and subject to scheduled purging of old files. User quota on the scratch storage is typically 100 TB and one million files.
- *Home storage*: The home storage was designed for use by individual researchers to hold files permanently. It is served from the DDN GS7KX filesystem appliance and uses a quota system for space management. Snapshots of user data are regularly taken and accessible from all Purdue research systems.
- *Archive storage*: The Fortress Archive (HPSS) system is a large, long-term, multi-tiered file caching and storage system utilizing both online disk and robotic tape drives. This archival system consists of 400 TB of disk cache, 1.2 TB of SSD for file metadata storage, and two tape libraries currently totaling 11 PB in storage capacity, which can be expanded to provide additional archival storage according to user demands.
- *Globus data transfer*: RCAC supports the use of Globus for transferring large amounts of research. Globus is used to reliably and securely transfer large amounts of data between systems. It also allows

researchers to easily share large datasets with collaborators through shared endpoints. Globus is available on all RCAC storage systems, including Research Data Depot and Fortress archive.

Visualization and Graphics

RCAC is home to the Purdue Envision Center, a facility that provides novel solutions to effectively communicate complex research concepts through computer graphics, advanced visualization, auditory (sound), haptic (touch), and multimodal interaction integrated with state-of-the-art high-performance computation. The Envision Center is staffed by a project manager, a technical lead, developers, and students to assist researchers, instructors, and leaders in their quest for new knowledge and innovative products. Recent and ongoing research supported by center staff includes projects in the life sciences, audiology, engineering, communications, theatre, computer science, geology, and management.

Science Gateways

Science gateways are online platforms that provide access to advanced resources for researchers, educators, and students. "Through streamlined, online, user-friendly interfaces, gateways combine a variety of cyberinfrastructure (CI) components in support of a community-specific set of tools, applications, and data collections."[1]

Purdue's RCAC provides two production-quality science gateways, *MyGeoHub* and *DiaGrid*, for use in research projects. Both gateways are built on the HUBzero™ Scientific Collaboration Platform open source software stack.[2] They provide "Science-as-a-Service," online applications delivered to user's web browser with graphical user interfaces while accessing large-scale computing and data resources behind the scene. On these open platforms, researchers and students can share their own scientific applications following a built-in, do-it-yourself process without having to be an expert in writing web applications. Publications of datasets and research materials are provided with Digital Object Identifiers (DOIs) to make them findable and accessible. These gateways provide many functions for collaboration and learning, including groups, projects, courses, wikis, discussion forums, reviews, tagging, and citations. They also support modern interactive programming environments, such as Jupyter Notebook and R Studio, that are accessible in a web browser. There is also a built-in ticket system for technical support. Both gateways are operational 24-7 and hosted and supported by the HUBzero team of professional staff. More specifically:

- *MyGeoHub* (http://mygeohub.org)[3] is a geospatial science gateway that specializes in support for geospatial data access, processing, analysis, and visualization. It was established to support integrative, geospatial, data-driven research activities and broad dissemination of the outcomes, and currently hosts more than 10 research projects and more than 10,000 annual users. MyGeoHub currently supports 360 accessible resources including courses and training materials, 55 tools, and more than 15,500 simulations run annually. Students and instructors from more than 30 U.S. and international higher education and research institutions have used resources on this site in their classes in the past 12 months.

  MyGeoHub features the new *geospatial data management and visualization capabilities* created by a recent NSF data infrastructure project named **GABBs** (Geospatial data Analysis Building Blocks), which further lowers the barrier for researchers to use, manage, and share geospatial data online. Release 2 of the GABBs software has been deployed on MyGeoHub featuring geospatial software libraries and widgets for building tools; built-in geospatial data management including viewing, data layer control, automatic metadata extraction, and geospatial search; data service API for linking tools

[1] Wilkins-Diehr, Nancy. "Special issue: science gateways—common community interfaces to grid resources." Concurrency and Computation: Practice and Experience 19, no. 6 (2007): 743-749.

[2] M. McLennan, R. Kennell, "HUBzero: A Platform for Dissemination and Collaboration in Computational Science and Engineering," Computing in Science and Engineering, 12(2), pp. 48-52, March/April, 2010

[3] Kalyanam, R., Zhao, L., Song, C., Biehl, L., Kearney, D., Kim, I. L., Shin, J., Villoria, N., & Merwade, V. (2019). MyGeoHub—A sustainable and evolving geospatial science gateway. Future Generation Computer Systems, 94, 820–832. https://doi.org/10.1016/j.future.2018.02.005

and data sets into workflows; and end-to-end tools that makes it easier for domain researchers and non-GIS experts to explore, analyze, and visualize geospatial data. Another project named **GeoEDF,** recently funded by NSF, will further extend the HUBzero platform with new capabilities that help seamlessly connect computing with large, remote, distributed geospatial datasets and build reusable workflow pipelines using the latest technologies such as Singularity containers, Kubernetes, and Jupyter Notebook. The GeoEDF software will be publicly accessible on this platform as it is released in 2020.

- *DiaGrid* (http://diagrid.org) is a computational science gateway focusing on simulation and other computational tools that need access to significant high-performance and high-throughput computing resources. Current applications on DiaGrid include tools for bioinformatics, molecular dynamics, structural biology, climate modeling, and statistical analysis and parameter sweeps using R language.

National Cyberinfrastructure Partner and Resource Provider

Purdue's RCAC is a current partner in the NSF-funded advanced cyberinfrastructure XSEDE (xsede.org). As XSEDE Extended Collaboration Service Support (ECSS) personnel, RCAC research scientists and software engineers provide in-depth advanced consulting services to help researchers nationwide to effectively utilize the XSEDE resources. Their expertise encompasses a variety of fields including code optimization, scientific application development, and science gateway development and operation. Purdue staff have created and led the successful Campus Champion program throughout the TeraGrid and XSEDE 1.0 projects. This program resulted in a grass-root community covering all campuses in the U.S. to help researchers access and utilize national advanced computing resources and share experiences and best practices. Purdue continues to coordinate program activities to work toward continued growth and sustainability.

Purdue's RCAC has been a resource provider to the national advanced cyberinfrastructure ecosystem through XSEDE (previously TeraGrid) and Open Science Grid. Its NSF-funded Anvil system, a new capacity computational resource, will be integrated with XSEDE and placed into production in 2021. Featuring the latest CPU processor, GPU, interconnect, and storage technologies, Anvil integrates a large-capacity high-performance computing (HPC) cluster with a comprehensive ecosystem of software, access interfaces, programming environments, and composable services to form a seamless environment able to support a broad range of current and future science and engineering applications. Open OnDemand and ThinLinc tools allow users to quickly become productive on Anvil through Linux and Windows desktops or familiar tools through their browser (e.g., Jupyter, RStudio). Complex scientific software environments and application stacks will be supported via containers orchestrated within a powerful composable subsystem. Anvil supports cloud-bursting of computational workloads as well as use of public cloud machine learning platforms including GPU and FPGA accelerators and software tools to automate hyperparameter tuning and algorithm selection for exploratory ML research.

Purdue is also a resource provider to the Open Science Grid with its Condor pool for high-throughput computing applications. Purdue provides computing and storage resources to the OSG to support the CMS (Compact Muon Solenoid) project as a Tier-2 site and other OSG virtual organizations.