

Personnel

Arizona State University (ASU) is served by both the central University Technology Office (UTO) and the Research Technology Office (RTO). UTO is the central IT organization with over 540 FTEs across multiple service areas including desktop support, wired and wireless networking, public and private cloud, identity management, information security, and web application development. UTO oversees core campus IT services such as payroll, email, instant messaging, user file storage, and document creation/collaboration. UTO also handles ASU policies regarding IT services, data governance, and information security. RTO focuses on IT services directly supporting research and researchers. Specifically, RTO comprises 65 FTEs covering Research Computing, Scientific Software Engineering, Research Data Management, Business Intelligence, and Web Services. RTO is overseen by the Chief Research Information Officer who reports to the University's Executive Vice President of Research. The Research Computing staff consists of computational scientists, programmers, engineers, and database administrators with expertise in all areas of computing, including scientific and parallel computing, big data analytics (in memory), custom software development, database engineering, and scientific visualization.

Advanced Computing

Research Computing is an academic supercomputing facility providing high performance computing (HPC) environments, a data intensive ecosystem, connectivity to the Internet2 research and education network, and large-scale data storage with elastic capacity to the public cloud. Research Computing provides a variety of HPC (both physical and virtual), cloud, storage, development, implementation, and consulting services. Research Computing consulting services and support for computational investigations, including data analysis, simulation, modeling, visualization, and other high-performance approaches include:

- Identifying optimal systems and software platforms
- Training in computational and/or graphics algorithms, tools, and packages
- Developing custom post-processing graphics tools
- Creating virtual environments for scientific research and fine arts
- Tuning applications for peak performance and implementing parallel algorithms and programs
- Purchase consultation for server, HPC, and storage acquisitions
- Virtual server provisioning (local, cloud)
- Physical and virtual server management
- Providing state-of-the-art interfaces to HPC systems
- Recharge of non-preemptable computing time and data storage solutions
- Accessing extensive external, government-funded compute resources (XSEDE, OSG)

Advanced Computing Systems

The Agave supercomputer is ASU's flagship high performance computing (HPC) cluster. Agave is a heterogenous Intel-based HPC cluster containing over 15,000 CPU cores. Each node is equipped with a solid-state drive and system memory ranging (depending on the node) from 128GB to 256GB of DDR4 2400 RAM. For large memory applications, the cluster also contains three nodes with 1TB, 1.5TB, and 2TB of DDR4 2666 RAM. GPU computing capabilities include access to over 330 NVIDIA A100, V100, K80, GTX 1080, and RTX 2080 GPU accelerators. A dedicated pool of 1.2PB high performance BeeGFS provides fast scratch storage for compute jobs. Compute nodes are accessible through four login nodes, one of which hosts the NSF-funded Open OnDemand web interface. Compute jobs are managed with the SchedMD Slurm scheduler. Agave is supported by a 100Gbps InfiniBand network fabric. It is connected to the campus Science DMZ, Internet2, and Data Transfer Nodes by a 100/40GE core network.

A dedicated pool of 1.2PB high performance BeeGFS fast scratch storage is presented to the Agave cluster via dual interconnected networks (InfiniBand and Omni-path), and a 1.8PB network attached storage array provides HPC home directory storage. For general purpose research data, a 4PB network attached storage array provides project storage.

Researchers may also purchase their own compute nodes and incorporate them into the Agave cluster, with Research Computing supplying all necessary rack space, power, cooling, networking, and software maintenance. Compute nodes are 52-CPU-core Intel servers manufactured and supported by Dell Technologies. Researchers and their delegates have priority access to their purchased nodes, and any idle capacity is made available to computing jobs for the general ASU research computing community. Such jobs will be guaranteed to run without preemption for at least four hours, after which jobs submitted by the node's owner will preempt them. The owner may also reserve their nodes exclusively for up to three one-week periods per year. Purchasing computing capacity in this manner allows researchers guaranteed access to the necessary computing power without needing to operate and maintain their own servers.

Research Computing will support researcher-purchased nodes for as long as feasible. However, beyond the hardware warranty period, the faculty is responsible for any hardware and labor costs necessary to maintain the hardware. Once the warranty period has expired, Research Computing may remove the node from the cluster if it is no longer technically feasible to support it.

Dell Center of Excellence for HPC and Artificial Intelligence

ASU has been designated a Dell Center of Excellence for HPC and Artificial Intelligence, the third such center in the United States, and the sixth globally. This distinction has grown out of a close collaboration with Dell HPC experts on system architecture, design, and innovation.

Through this new partnership ASU is deploying a new high performance computing system in the second half of 2021. Anticipated characteristics are a mixed CPU/GPU environment of approximately 15,000 CPU cores, 200 GPU devices, and 6PB of high performance storage.

Open Science Grid

Research Computing runs a 20-node Open Science Grid (OSG) site for the research community at large and is investigating using spare cycles on idle lab workstations to significantly augment this OSG infrastructure.

Data Center

ASU Research Computing maintains an on-campus data center in Interdisciplinary Science and Technology Building 1 (ISTB1), in the center of the ASU Tempe campus. The ISTB1 facility was built in 2012 consists of a 5,000 square foot primary data center for critical systems, networking, and computational resources as well as a 3,000 square foot secondary data center for non-critical systems, development, and individual research development equipment. Both data centers employ a standard “hot and cool aisle” layout with computer room air conditioning units totaling 200 tons, supported by campus chilled water systems. Room-dedicated FM-200 fire suppression systems protect the facility. Power for the facility is from on-campus natural gas turbines fed by Utility natural gas. Data center power supports dual power feeds, protected by two uninterruptible power supply units totaling 1MW, with an onsite diesel generator providing 2MW of power and a 500-ton emergency chiller in the event of a loss of utility power. Access to the data center requires a keycard and PIN, and the facility is monitored 24x7 from a

dedicated operations center, and physical access is controlled and maintained by the UTO operations center.

ASU Research Computing is currently building out a new data center at the Iron Mountain Phoenix facility. This facility will provide ASU with more than four times the capacity of the existing infrastructure in a commercial Tier III+ data center with advanced power, cooling, and network capabilities. ASU Research Computing will be a core tenant of the facility and will have the capabilities to provide secure research (up to FISMA¹ High) security. A private fiber ring will connect the facility to the ASU Tempe Campus. Internet2 and Commodity Internet circuits will be available at the facility as secondary connectivity, as well as private point-to-point circuits as Research requires.

Network

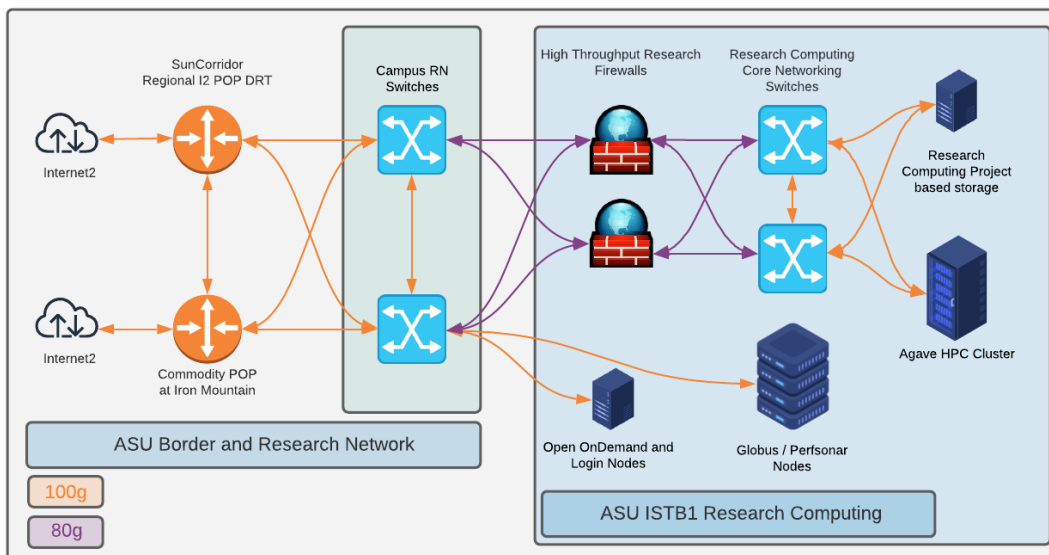


Figure 1: Research Computing Network Connectivity

Figure 1 illustrates the current ASU network border topology. Primary network access is via a 100 Gb Internet2 circuit to the Tempe campus. Secondary 10 Gb commodity Internet circuits provide additional and partially redundant network access directly from the Tempe campus. The ASU Polytechnic campus (25 miles southeast of Tempe), the ASU West campus (25 miles northwest of Tempe), and the ASU Downtown campus (10 miles west of Tempe) connect to the Tempe campus via redundant 10Gb circuits on a commodity fiber ring. The ASU network is monitored 24x7x365 by a commercial network provider as well as by the University Technology Office.

Buildings on all ASU campuses are connected in a hub-and-spoke model, with most buildings served by redundant 10Gb links and 1Gb to end users. The campus network employs an advanced security complex consisting of a layered defense-in-depth deployment of security controls that include DDoS and IP reputation, a variety of specialized network firewalls, and anti-phishing protections. The ASU cybersecurity program also includes mandatory security education and awareness training, and the UTO

¹ <https://www.dhs.gov/cisa/federal-information-security-modernization-act>

Governance, Risk, and Compliance Team conducts continuous assessments evaluating risk and vulnerabilities.

Science DMZ

The ASU Science DMZ is a network enclave that bypasses the network security complex. The Science DMZ is explicitly designed for high-throughput data movement, incorporating 100/40 Gigabit Ethernet, virtual circuits, and software-defined networking capabilities as well as dedicated systems for large data movement requiring a friction-free path, with security policies and enforcement mechanisms tailored for high performance science environments.

Data Storage

The ASU University Technology Office (UTO) supports cloud storage using a variety of cloud-based storage offerings, including Enterprise Dropbox (for Staff/Faculty, 1TB limit), Microsoft OneDrive (available to all Staff, Faculty, and Students via Office 365, 1TB limit), and Google Drive (available to all Staff, Faculty, and Students via G Suite for Education, no storage limit). UTO provides storage to Business Units via SMB/CIFS on Enterprise NetApp Network Appliances.

ASU Research Computing provides 100GB of home directory storage for users of the Agave Cluster, as well as access to the 1.3PB BeeGFS high-speed short-duration scratch environment for cluster computing jobs. Research Computing also operates 4PB of network-attached project term storage. This storage is accessible to the Agave HPC cluster and individual researcher workstations via traditional network shares. The Globus data movement platform provides resilient high-speed access to data stored on Research Computing systems and allows for transfer to user's University provided Google Drive accounts.