

# Syracuse University Facilities, Equipment, and Other Resources

## Green Data Center

In the fall of 2010, Syracuse University completed construction of the [Syracuse University Green Data Center \(GDC\)](#). In partnership with IBM and the New York State Energy Research and Development Authority (NYSERDA), 6,000 square feet of datacenter space was added that is powered by a unique tri-generation power plant which utilizes water chilled racks to increase power and cooling efficiencies. On site natural gas micro-turbines provide 650 kilowatts of power and their waste heat is used by two Thermax Absorption Chillers to provide equipment and building cooling. The GDC cooling capacity is approximately three times that needed by the data center; excess chilled water is used to provide air conditioning for adjacent campus buildings. Lead-acid batteries provide 17 minutes of emergency backup power (at full capacity) in the unlikely event that the turbines and the utility grid simultaneously fail. To help further research on green data center practices the GDC has been instrumented with hundreds of sensors. It is now also the main center for production computing resources, marrying research and administrative computing interests. The internal GDC network is 10 Gigabit with redundant connections to all hosts and network components whenever possible. The GDC supports a robust virtual private cloud in which 95% of campus servers have been consolidated. This includes servers (both research and operational) on campus that were housed in poor environmental conditions within distributed locations.

The GDC is connected to the campus network and the original data center via two bundles of geographically diverse, 144 strand fiber paths. This can be and has been used to provide direct connectivity between a researcher's campus building and the hosted area in the GDC.

## Green Data Center Hosting for Researchers

Roughly half of the space in the GDC is hosted space designed to provide a secure physical environment that is flexible enough to allow access to the equipment by researchers and graduate assistants. The hosted area is caged, and provides a separate entrance that allows access through a combination of ID cards, biometric finger prints, and PIN numbers. Over half of the computing load in the GDC serves research computing needs. Researchers on campus are able to rely on having space in the datacenter that is very low cost and provides ideal environmental conditions for their equipment with power redundancy from the multiple layers of protection including the traditional electric grid, natural gas fired turbines (with on-site propane gas storage), and UPS. Work continues to encourage researchers who host equipment under less than optimal conditions in their departmental buildings to migrate their approved equipment into the GDC.

## OrangeGrid

This distributed computing system, comprising some 12,000 cores, is used by SU faculty and researchers, particularly in the physical sciences and engineering, who need reliable, high throughput computing (HTC). The computers in the grid are optimized to perform a large number of smaller parallel jobs (typically less than 24 hours), providing high processing capacity over long periods of time. The grid utilizes virtualization via Oracle's VirtualBox, scheduling via the HTCondor HTC System, and the HTCondor Virtual Machine Coordinator (CVMC), a small application developed by SU's Information Technology and Services (ITS) team to manage the multiple desktop components. These components are

distributed to desktop clients via Microsoft's Active Directory. HTCondor, developed with support from the National Science Foundation, manages the grid's workload. The computer's task scheduler detects when its host computer is idle, starts up CVMC, and connects to HTCondor to receive work. When user activity is detected on the computer, research operations are immediately stopped. The use of virtualization acts as a barrier which separates the researcher and their content from the user's information on the same computer.

### Academic Virtual Hosting Environment (AVHE)

The AVHE provides a private compute cloud to the Syracuse University research community, lowering the bar of entry for small to medium sized research efforts. The private cloud uses virtualization to provide flexibility and hardware sharing to allow multiple researchers to operate on an underlying server and storage infrastructure. This lower bar of entry and flexibility provides an environment that supports both traditional and non-traditional computational research. Another use case for the AVHE is the building of small to medium sized clustered research computing environments, reducing the need for researchers to build and maintain small physical clusters. Services within the AVHE include medium scale storage for research data storage and archival. To date this includes over a half of a Petabyte of research data and output. The AVHE utilizes the abilities of virtualization to provide high availability, which automatically migrates workload to alternate resources in the case of physical server failure. Backup services within the AVHE are a standard service provided to researchers.

### Crush

Crush is a medium scale virtualized research cloud. It is designed to be allocated for compute intensive work and to be used in tandem with the AVHE where the data and work scheduling infrastructure are maintained. A primary use case for Crush is the provisioning of a “cluster within a cluster” providing access to dedicated, customized compute nodes for high-performance and high-throughput computing. Currently Crush consists of 10,000 cores and 45 Terabytes of memory.