Florida State University
School of Computational Science

Integrated Technology Plan

Fiscal Year 2007-2008

CONTACT INFORMATION AND ORGANIZATION

School of Computational Sciences (SCS)

Director of SCS:

Dr. Max Gunzburger (Secretary: Cecelia Farmer, 644-0143, farmer@scs.fsu.edu)

voice.....644-7060 e-mail..... gunzburg@scs.fsu.edu

Associate Director of the SCS, for Computing and Director of the Shared-HPC:

Dr. James Wilgenbusch

voice.....645-0307 e-mail..... jwilgenb@scs.fsu.edu

Financial and Administrative Issues:

Michele Locke

voice.....644-2273 e-mail..... michele@scs.fsu.edu

SCS Local Systems Committee:

- James Wilgenbusch (Chair), Research Associate, School of Computational Sciences
- Kyle Gallivan, Professor, Department of Computer Science
- Wei Yan, Assistant Professor, Department of Chemistry and Biochemistry

FSU Shared-HPC Advisory Panel:

- Eric Chassignet (Chair), *Professor, Meteorology/COAPS*
  - E-mail: echassignet@coaps.fsu.edu
- Robert van Engelen, *Computer Science/SCS*
- Kyle Gallivan, *Professor, Computer Science/SCS*
- Hugh Nymeyer, *Assistant Professor, Chemistry and Biochemistry/SCS*
- Peter Beerli, *Assistant Professor, Biology/SCS*
- Sachin Shanbhag, *Assistant Professor, Chemical and Biochemical Engineering/SCS*
- Paul Eugenio, *Assistant Professor, Physics*
- Scott Steppan, *Associate Professor, Biology*
- Steve Morey, *Assistant Scholar Scientist, COAPS*
- Phil Cunningham, *Associate Professor, Meteorology*
- Bob Hart, *Assistant Professor, Meteorology*
- Steve Cocke, *Associate Scholar Scientist, COAPS*

Technical Support Staff SCS:

- Tom Green, *MS Computer Science*
- Jazcek Braden, *Ph.D. Physics*
- Xiaoguang Li, *Ph.D. Mathematics*
- Chris Cprek, *B.S. Computer Science*

Technical Support Staff HPC:

- Jeffery McDonald, *Ph.D. Physics*
- Paul Van der Mark, *Ph.D. Computer Science*
- Dan Voss, *MS Information Sciences*

Organizational Chart:

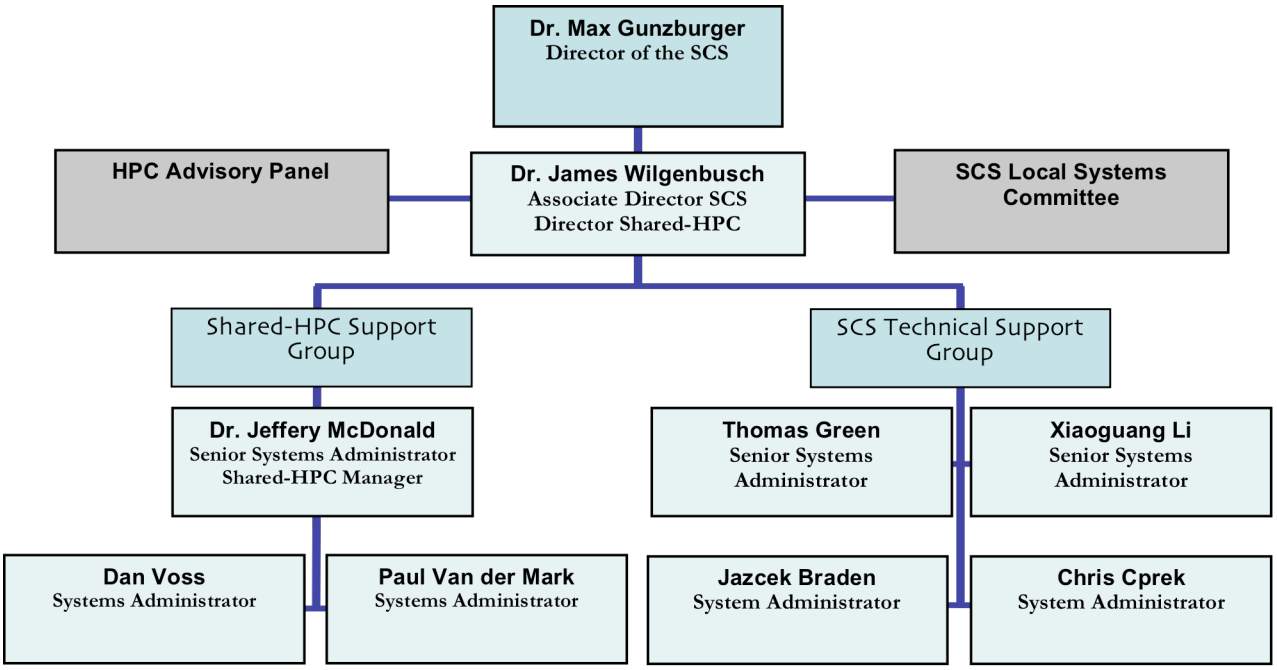

COMPUTATIONAL INFRASTRUCTURE

The SCS plays a major role in the support of FSU's cyber infrastructure by providing facilities and technical expertise in the support of scientific computing. The SCS manages two leading edge computing facilities for FSU; the SCS computing facility (located on the main FSU campus in Dirac Science Library) and the FSU shared-HPC facility (located at Innovation Park in the Sliger Building). Both facilities foster nationally recognized research programs, which span multiple disciplines and have the potential to attract more external funding. The SCS facility provides a highly flexible computing environment designed to support specialized and experimental hardware and software systems; while the shared-HPC supports a relatively homogenous and highly optimized production-based computing system.

SCS Computing Facility

The SCS facility supports an aggregate of over 1000 CPUs representing an assortment of computer architectures, interconnects, and operating systems. Systems hosted in the SCS facility are owned by SCS and non-SCS research groups and are dedicated to a wide range of research problems including; molecular biophysics, evolutionary biology, network modeling, and Monte-Carlo algorithm development (see http://www.scs.fsu.edu/twiki/bin/view/Computing/WebHome for more details). Hardware hosted in the SCS machine room is purchased by university and external funds. The SCS facility is equipped with two 40- ton HVAC cooling units, 1000 ft^2 of raised floor, an extensive power distribution system, and two large-format PowerWare UPS battery backup systems. Plans are currently underway to install a diesel-powered backup generator, which will provide backup power to all of the hardware and HVACs in this room. The SCS network is built on a 10 Gbps Cisco backbone, providing connectivity to a switching infrastructure and to key servers and storage. The SCS network connects to the 10 Gbps FSU campus backbone via 1 and 10 Gbps uplinks, which in turn connect to the 10 Gbps Florida LamdaRail.

FSU Shared-HPC

The SCS also manages the FSU shared-HPC facility, which supports multidisciplinary computing for the FSU research community. The shared-HPC is in the first phase of construction and

currently provides 2.3 TFLOPS of throughput. Phase II, slated for summer of 2008, will likely increase the throughput of the HPC by a factor of three. The HPC system consists of 128 Dell PowerEdge SC1435 compute nodes (512 CPUs) and 4 Dell PowerEdge 6950 head nodes. A Cisco 288-port non-blocking DDR infiniband switch and a 1 Gbps Cisco 6500 Ethernet switch for management and storage connected servers. All compute and log in nodes have access to a 78 TB Panasas high performance parallel Object Storage Devise. The shared-HPC is housed at the FSU central computing facility, which has a robust environmental and electrical infrastructure to accommodate current and future computing requirements. The HPC network infrastructure is also connected to the 10 Gbps Florida Lambda Rail and FSU Campus backbone.

Scientific Visualization

The SCS supports a general access laboratory for scientific visualization on the fourth floor of the Dirac Science Library. The Visualization Laboratory is equipped with four specialized workstations, three with GeForce 7900 video cards and the fourth with a Quadro FX4500 and a GeForce 8800 video cards. Work is currently underway in the Lab to use the CUDA SDK to better utilize the GPUs on these machines for general scientific computing. Each workstation has 4 GB of RAM as well as a wide range of specialized software for scientific visualization. The workstations have access to over 15 TB of high performance storage.

Thanks to a Major Research Instrumentation grant from the National Science Foundation, the SCS seminar room, adjacent to the Visualization Lab., was recently (December 2007) equipped with a high-resolution stereographic projection system to support multidisciplinary scientific visualization. Four state of the art rear-mounted projectors illuminate an 18' x 8' screen. The system switches between 2d and 3d mode with a simple touch of a button and also support numerous other input devises (e.g., a document camera, DVD/VHS player, cable TV, and two hookups for personal laptops) via a simple to use touch panel screen.

General SCS Infrastructure

The SCS manages approximately 30 servers for core network services using primarily generic Intel-compatible servers running CENTOS (a free version of Red Hat Enterprise Linux). This infrastructure is also used to support the shared-HPC and includes the following services.

Electronic Mail, utilizing IMAP4 and POP3 standards, with SSL security and SMTP authentication, including a web mail client, and incorporating spam and virus filtering.

List Server platform using the MailMan package services over 50 SCS and HPC sponsored lists.

Condor Server platform is used as a central job status collector for our distributed computational infrastructure. This platform matches jobs submitted on a central server to owner-based clusters and desktop workstations so that these systems can be used for scientific computations when they are not in use by their primary owner. This platform is associated with a large disk array that serves as a checkpoint image server for jobs submitted to the Condor cluster.

Web Server platform supports the SCS and HPC web sites as well as personal and course web sites for SCS faculty and students.

File Server platform hosts the home and research directories of SCS faculty and students. The system currently supports over 25TB of user and research space in a RAID5 configuration, plus over 25TB of scratch space. The system runs on a leading edge high performance parallel Object Storage Devise (OSD) designed for high-availability and high performance.

Backup of user home directories and e-mail is conducted nightly using a multiple 10 TB i-SCSI storage devise configured in Raid 5 with a hot spare.

Version Control Systems, including SVN and CVS implementations, support SCS students, faculty, and their collaborators.

Calendar service is available using the open source Kronolith calendar application. This system provides an integrated way for research and administrative groups to schedule meetings and to coordinate facility usage.

Trouble Ticket System is available using the open source system called the "Open Ticket Request System." This system gives SCS and HPC users a simple way to report potential problems with the SCS or HPC systems and gives our systems administrators a way to track the status of support requests.

Domain Name Service and DHCP are setup in a redundant manor to provide maximum availability. The SCS manages a large network space with numerous machines in the public IP space in order to support GRID and distributed computing requirements. An enterprise class Cisco Firewall, ACLs on key switches, and a number of client-based systems are used to protect the SCS systems from network based attacks.

Authentication services are provided via SSL encrypted LDAP. Three LDAP servers are configured in a redundant manor to ensure high available and to protect against data corruption.

PLANNED PROJECTS

Access Grid and Video Conferencing

The recently renovated SCS seminar room will be used to host a general access video conferencing solution that will be available for FSU faculty through a reservation system. The solution will build on the projection system and AV components recently installed in this room. The cost of the planned system will be relative low because the projection system, speakers, and echo cancellation devises have already been acquired and installed. Funds for relatively inexpensive video cameras, microphones, and staff time have already been allocated for this project. The Access Grid (http://www.accessgrid.org) will be used to provide the large format interactive multimedia video conferencing. The Access Grid is an ensemble of resources including multimedia large-format displays, presentation and interactive environments, and interfaces to Grid middleware and to visualization environments. These resources are used to support group-to-group interactions across the Grid. The Access Grid is used by well over 500 academic, research, and commercial organization throughout the world and will provide an excellent and cost effect way to foster and strengthen collaborations. The Access Grid also provides numerous educational venues, which will be used to support the research and educational objectives of the University.

Grid and Distributed Computing

The SCS uses the batch system known as Condor for queuing, scheduling, and prioritizing compute-intensive jobs. The Condor system is developed by the Condor team at the University of Wisconsin. Condor matches user submitted jobs to available computer resources, such as, clusters, servers and desktop workstations. As long as the machine has the required resources and does not interfere with the primary owners process, jobs will be sent to the machine and run. The SCS has effectively used this system for over six years to meet the computing needs of numerous SCS and affiliated faculty. The SCS plans to expand this system for opportunistic computing to include other computer laboratories and HPC resources on the FSU campus. The project is currently being lead by Jazcek Braden and Tom Green with help from SCS graduate students. Hardware resources for this project have already been allocated and deployed.

EMERGENCY PREPAREDNESS

Many of the computer systems and services maintained and operated by the TSG are critical to the research, education, and outreach components of the SCS mission. Failure of any one of these critical computer systems or services can have a wide range of effects, from the loss of email access to the loss of computer cluster functionality. To be prepared for disasters we first identify the computer services that are critical to the SCS mission and then specifically define the way in which members of the TSG are expected to monitor and to respond to the failure of a critical service.

Critical Services:

The following is a list of the services that are critical to some component of the SCS mission. Included implicitly in this list are the underlying systems and related services required by the services listed here to function properly.

- Mail (e.g., imap, pop, smtp, IMP)
- Web servers (SCS, HPC, and Morphbank)
- The principal SCS shared file system
- Network services (e.g., DNS, DHCP, and Firewall)
- Version Control Systems (e.g., CVS and SVN)
- Condor
- LDAP

Monitoring and Maintenance of Critical Services:

The SCS technical support group (TSG) is responsible for monitoring the mission critical services listed above. Other services are also monitor, however only the above warrant a critical response. A designated member of the TSG is on call 24 hours a day and seven days a week. In the event of a critical service failure, notification is sent within 30 minutes of the failure to a pager and to an email address. The on call person is expected to initiate restoration procedures for the failed service within one hour of the time that a service failure notification is sent. On call duties rotate among qualified members of the TSG on a weekly basis. The following members of the TSG share on call duties: *Tom Green, Xiaoguang Li, Chris Cprek,* and *Jazcek Braden*.

Data Protection and Backup

SCS provides a massively parallel data storage solution using an object-based system called Panasas. The Panasas system offers a number of features to minimize downtime and to maintain a reliable, robust, and highly available storage solution. These features include;

- The scope of failures are limited to the object itself, and do not impact the whole disk
- Object-based rebuilds take seconds rather than hours
- Panasas provides "vertical parity", which means larger disks are no more failure prone than smaller disks
- Storage is configured in RAID 5 to limit data loss in the event of a disk failure

In addition to the inherent protection provided by the Panasas storage system, daily backups to two iSCSI disk arrays in RAID 5 configuration are also made in a 28-day rotation. Files on individual desktop systems throughout the SCS are not automatically backed up. Users can save files to the shared file system by mounting the shared file system locally and copying important files to their home directory. The SCS backups are performed using a simple but robust system based on the rsync program.

Outreach: Support for non-SCS Computing Systems

Eligibility

All units within the College of Arts and Sciences are eligible to have machines hosted in the SCS computer room. The SCS's Local Systems Committee will evaluate each request on a case-by-case basis and will agree to host machines provided the capacity allocated for non-SCS systems has not been reached and the requirements listed below are met.

Requirements

Each of the requirements enumerated below is described in greater detail in the following pages.

1. The SCS is not responsible for incidental costs associated with the installation or maintenance of computer systems hosted in the SCS computer room.
2. Requests to house computer systems in the SCS computer room must include a detailed hardware proposal.
3. Requests to house computer systems in the SCS computer room must include a proposal describing how the computer systems will be managed.
4. Computer systems hosted in SCS computer room are required to run the SCS load distribution system.
5. Requests to house computer systems in the SCS computer room must include a designated faculty liaison from within the unit making the request.

1. Incidental costs

In addition to the substantial costs of providing a cold room with reliable power, there are a number of "incidental" costs associated with hosting computer systems. These costs include server racks, power distribution units (PDUs or power strips), and network switches. The SCS is not responsible for covering incidental costs. At a minimum, the SCS will provide a single outlet for each electrical circuit required to power the computer system and a single 10/100/1000 network uplink in the vicinity of the computer system. The cost of incidentals (e.g., racks, PDUs, and network connectivity) can either be included in the price of a fully integrated system or the SCS will provide these items at cost. In order to avoid half empty racks, the later option may be

required. For example, if a group plans to purchase a small number of servers, then the SCS would prefer to use an existing rack, PDU, and network switch. In this case, costs will be assessed according to the portion of the resource being used.

2. Hardware proposal

Requests to house computer systems in the SCS computer room must include a detailed hardware proposal. An agreement to house a computer system in the SCS computer room is limited to the hardware described in the proposal. This requirement is intended to safeguard against someone purchasing hardware that is in some way incompatible with the SCS machine room infrastructure. In addition, the hardware proposal is needed to estimate the power, space, and cooling requirements of the computer system so that we do not exceed the capacity of the SCS computer room. The SCS has a wealth of experience with numerous hardware vendors and can create the hardware proposal on the behalf of the unit making the request.

Hardware proposals must include the following information:

- Server rack dimensions*
- Switch type and size*
- PDU type and quantity*
- Chassis type per unit
- CPU type and quantity per unit
- Motherboard type per unit
- Ethernet or network adapter type and quantity per unit
- Hard drive type and quantity per unit
- Memory type and quantity per unit
- Type and quantity of all peripheral devises that will be part of the computer system (e.g. disk arrays, CD/DVD drives, etc.)

* if not using equipment provided by the SCS

3. Management proposal

Requests to house computer systems in the SCS computer room must include a proposal describing how the computer system will be managed. Two general administrative options are available; 1. the SCS will manage the computer system, and 2. the unit making the request will manage the computer system.

Option 1: SCS manages the computer system

Option 1 gives a group access to a computational resource without the hassle of having to administer it. The SCS will agree to manage a computer system only if the SCS is given discretion to choose the hardware, the operating system, and the tools used to update and distribute software on the

computer system. The SCS systems manager determines who is granted access to privileged commands. Access to privileged commands will be controlled by the "sudo" command; root logins are not permitted. Only members of the SCS systems group will know the root password. The SCS systems manager will determine the update cycle of the core operating system and system software.

Option 2: Non-SCS unit manages the computer system

If a unit chooses to manage their own computer system, then a qualified systems administrator must be designated to manage the resource. The required management proposal must address in detail how the designated systems administrator will tend the following tasks: account management, disk management, network management, system installation, software management, security, systems management, user support, hardware support, and administration of purchase and warranty information. Systems not managed by the SCS will be relegated to a separate secure network. A firewall will monitor incoming and outgoing network traffic and the SCS systems manager will determine the stringency of the firewall. Access to the machine room will be limited to the designated systems administrator. Members of the SCS systems group will have root access to all systems housed in the SCS computer room.

4. Cycle sharing

In order to maximize throughput on SCS and non-SCS computational servers and cluster nodes, the SCS uses a job scheduling and queuing system called Condor (http://www.cs.wisc.edu/condor/). All of the computer systems housed in the SCS computer room must participate in this system. The Condor system allows fine-grain configuration of policies to support distributively owned computing resources. For example, specific nodes can be configured to "prefer" jobs submitted by specific users. This way, a resource owner will have limited access to all of the SCS computational resources and guaranteed access to his or her machines. This system ensures that cycles are not wasted if the primary resource owner is not using them.

While the Condor system provides some important advantages to managing and distributing computationally intensive jobs, a resource owner is not obligated to use it. For example, a resource owner can run jobs interactively on any of his or her computer systems and his or her jobs will preempt jobs managed by the Condor system. Circumventing the Condor system in this way is to the disadvantage of the resource owner, however. By submitting jobs through the Condor system, a resource owner will not only have guaranteed access to his or her machines, but the Condor system will also provide access to the other machines housed in the SCS machine room.

5. Faculty Liaison

Requests to house computer systems in the SCS computer room must include a designated faculty liaison from within the unit making the request. The liaison will be responsible for conveying information to the users of the computer system housed in the SCS computer room. If the SCS agrees to manage a computer system, the faculty liaison is responsible for approving account requests, requesting changes to a user's priority, and requesting software installations. The authority to make these requests can be transferred to a designate person (e.g., graduate student) at the written request of the faculty liaison.