Center for High-Performance Computing @ the University of Utah

News

University of Utah Fall 2006

Article

Computational Science and Engineering Online (CSE-Online): A Laboratory without Walls for Scientific Computing

by Prof. Thanh N. Truong
Department of Chemistry, University of Utah

> *…a new age has dawned in scientific and engineering research, pushed by continuing progress in computing, information, and communication technology, and pulled by the expanding complexity, scope, and scale of today's challenges. The capacity of this technology has crossed thresholds that now make possible a comprehensive 'cyberinfrastructure' on which to build new types of scientific and engineering knowledge environments and organizations and to pursue research in new ways and with increased efficacy. —NSF Blue Ribbon Advisory Panel report on cyberinfrastructure, 2003*

In accord with the above principal finding of the NSF Blue Ribbon Report, enabling research to be performed in new ways and with increased efficacy is the aim of the Computational Science and Engineering Online (CSE-Online).

In order to understand how CSE-Online accomplishes its aim, it is important to recognize that current research environments in a variety of computational science domains require an understanding of the computational methods, technical skills in operating computers (detailed knowledge of UNIX, parallel processing and job submission environments, and file management and exchange), and the ability to edit, convert and format the necessary files. The researcher also has to learn each of the requisite tools, including the different look and feel for each and different means of interacting and running each program. Data from remote web servers must be first downloaded to the user's local desktop computer before using them as input data for other applications. The researcher needs to be familiar with the input/output file formats and execution of all application tools for all steps necessary to complete the project. They need to transfer data and remotely login to specific computers to submit jobs. This requires learning the vagaries of the different operating and job submission systems. Filtering and parsing data from one application to another has to be done manually. Communication with collaborators has to be done conventionally, i.e. by phone, fax, and/or email.

The CSE-Online vision is to make all computing resources available to the user at their fingertips by creating an innovative virtual desktop environment that allows data,

CHPC AT SC2006

CHPC, in conjunction with the Scientific Computing and Imaging Institute (SCI) and the Center for the Simulation of Accidental Fires and Explosions (C-SAFE), will again have a booth at this year's ACM/IEEE SC2006 conference in Tampa, FL. This year we have taken our collaborative efforts to the next level by combining forces with the newly formed Utah State University Center for High Performance Computing (HPC@USU). This multi-school presence is an effort to better showcase high performance computing and networking in the state of Utah and our continuing effort to create permanent means for collaboration between the two institutions.

Our joint booth will highlight work being done on both campuses aided by high performance computing systems. Posters, movies, images and stereo 3-D visualizations will be used to showcase the diversity of supported projects. Brian Haymore, our Lead Computation Cluster Administrator, will also be giving live demonstrations on CHPC's NFSroot environment.

CHPC presence at the conference will extend outside the booth:

POSTER SESSION: "DIGITAL SHERPA"
"Digital Sherpa" is a grid application for executing jobs on many separate HPC resources at a time. The Digital Sherpa poster will be presented by CHPC's Ron Price, Victor Bazterra, Wayne Bradford and Julio Facelli.

SC GLOBAL SHOWCASE: "INTERPLAY: DANCING ON THE BANKS OF PACKET CREEK"
Wednesday Nov. 15, at 1:30 p.m., Room 14 - 16. This performance directed by Jimmy Miklavcic will be a telematic event, utilizing Access Grid technology to bring together five sites across the country into an exciting collaborative performance work.

SC GLOBAL EVENT: "PERFORMANCE AND PRESENTATION PRODUCTION ELEMENTS"
Wednesday Nov. 15, at 2:30 p.m., Room 14 - 16. Beth and Jimmy Miklavcic will discuss and demonstrate techniques that go into creating effective video and audio streams for both Access Grid meetings and performances

applications, and computing resources to be delivered to the user's desktop transparently. By doing so, it introduces a paradigm shift to research and education in computational science and engineering.

We define the computing resources available to the user as consisting of all resources (data, application tools, and computing cycles) including those on the user's desktop. For data, this includes private data on their desktop computer and on all remote servers for which the users have login accounts, public data at public databases available online, and possibly shared data which are made available from others. For application tools, this includes the user's own tools, CSE-Online tools, other third-party open-source or commercialized tools. For computing resources, this includes the user's desktop CPU, remote servers on which the user has CPU allocations, and compute grids such as the TeraGrid.

How CSE-Online accesses application tools

Upon creating a workspace for the first time, the user fills up their CSE-Online desktop with a number of icons of built-in Java applications chosen from the catalog. These icons contain the knowledge about the application, namely the location of the executable (jar file), how to initiate it, its version, etc. Double clicking the icon of the application initiates a sequence of processes, particularly:

1. Is there a version of the application in local cache? If not, download the executable from its server and run it.
2. If yes, then check if the version on the local computer is current. If no, update the local one by downloading the new version, then run it.
3. If the local version is up to date, then run it.

This way, the user always has the latest published version of the application and the developers do not have to worry about sending out bug fixed versions. This gives CSE-Online its online and on demand characteristics.

The CSE-Online software framework allows the applications to be independent from each other. Thus, they can be developed independently. They can communicate with each other via open application programming interfaces. It allows users to bring their own tools as well as the third-party tools into the desktop environment and allows these applications to communicate with other tools in the environment. This makes CSE-Online extendable and its applications interoperable and modular.

The above is true for Java applications. What about server-based applications? In fact, 99% of software for computational science are server-based due to their computing intensive nature and many are commercialized and thus have license restrictions. How does CSE-Online fulfill such license restrictions?

CSE-Online currently supports a number of server-based applications such as Gaussian03, GAMESS, AMBER, TheRate, Thermo, etc. These applications have different license restrictions. To access these tools CSE-Online has Java Graphic User Interfaces (GUIs) to create input files. Depending on the license restrictions, CSE-Online will invoke different mechanisms for initiating the server-based applications with the input files from the GUIs. For instance, TheRate is a CSE-Online server-based tool for calculating rate constants of elementary reactions. It does not require significant computing resources and thus we provide it as a free service to the community. In this case, CSE-Online allows the user to request the service, i.e. running TheRate on the CSE-Online server, using Web-Services technology. The users do not have to install TheRate anywhere. On the other hand, Gaussian03 is a commercialized software for quantum chemistry which has site licence restrictions.

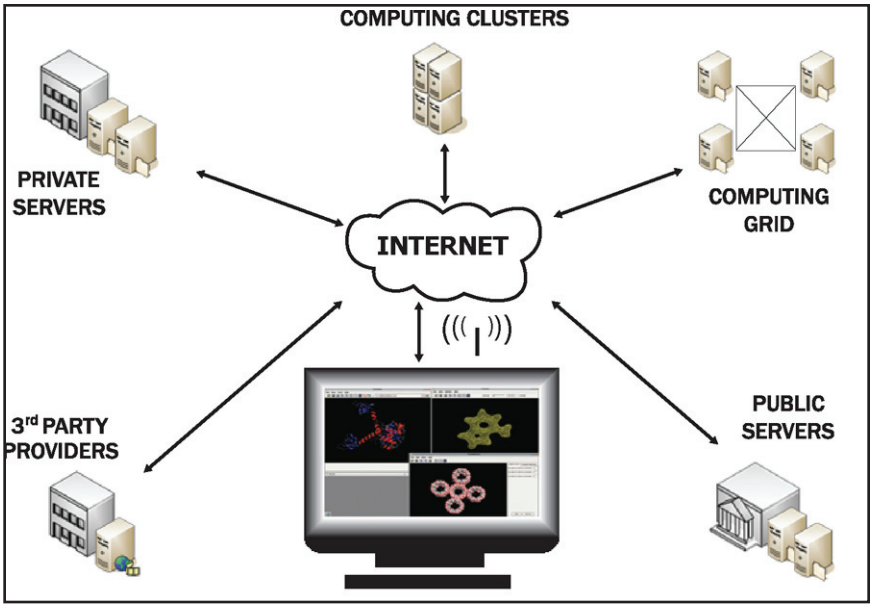

Did you know?

CHPC has a large memory node available by request. This node on landscapearch has 4 - 2.4 Ghz processors and 32 Gbytes memory. If you can demonstrate that you could make significant scientific progress by gaining access to this node, please send an email request to *problems@chpc.utah.edu* describing the computational details of how your research would benefit.

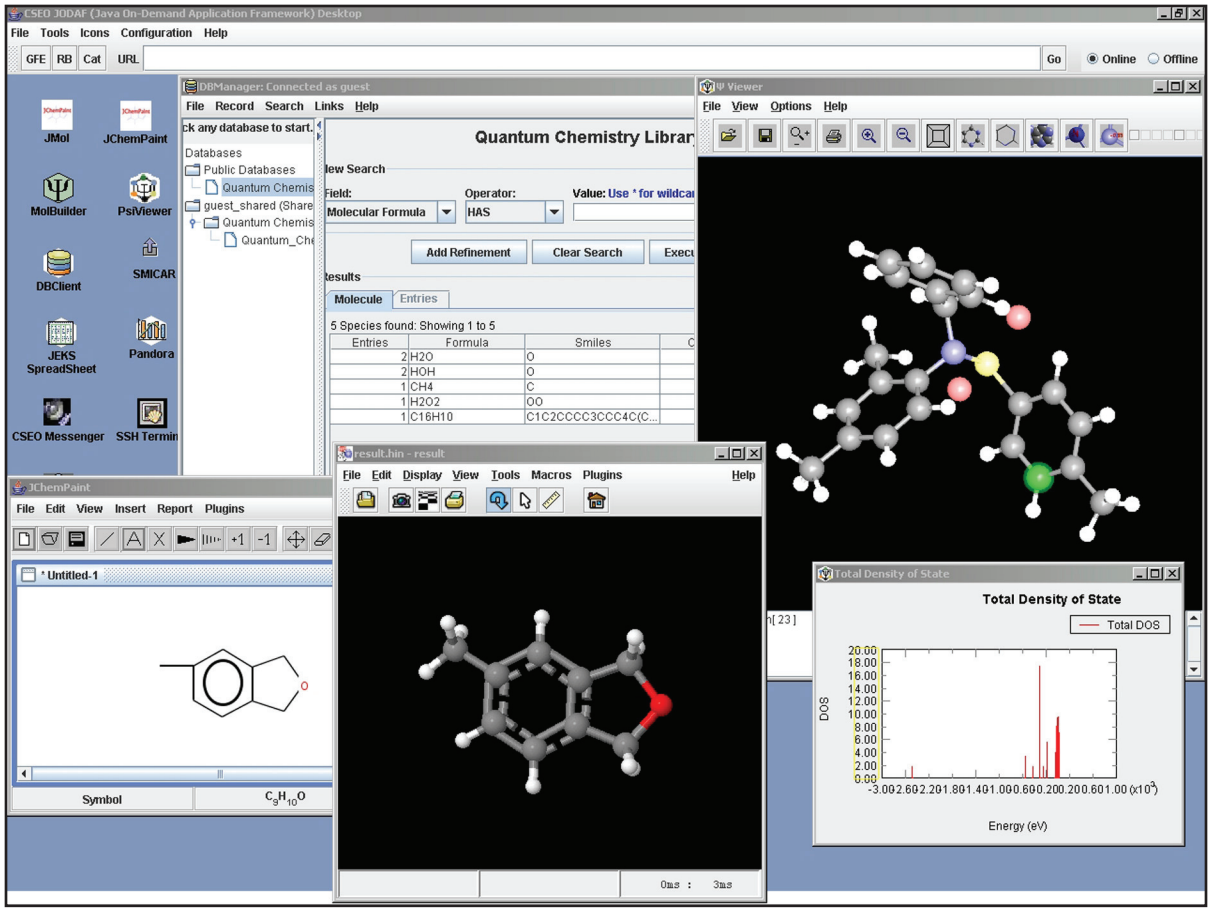

The Gaussian03 GUI creates the input file then sends it to the Resource Broker tool for submission. The users must configure the Resource Broker for the location of the Gaussian03 executable on the server on which they plan to run the job and the Resource Broker will submit the job if the users are authorized to use the server. This way CSE-Online can fulfill license restrictions of any server-based application.

Users can also develop Java GUIs to integrate their own tools into CSE-Online to take advantage of the Resource Broker for submitting and monitoring jobs.

How CSE-Online accesses data

There are two kinds of data depending on their sources: files from file servers and records from databases

For files, CSE-Online has an application called Global File Explorer (GFE), that allows users access to files from their local desktops as well as all remote accessible file servers. Users can copy/move files between servers by drag-and-drop. GFE simplifies the task of managing files on multiple servers. GFE also allows users to share files with others (not using the same servers) easily. This share feature is currently being implemented. For a record from databases, similar to GFE, CSE-Online provides a database client (DB Client) environment for accessing databases. DB Client distinguishes three types of databases depending on their accessibility: public; private; and shared databases.

There are two types of public databases: those hosted by CSE-Online and those hosted by third-party providers. CSE-Online hosts a number of public databases which are available to all users such as the Protein Data Bank DB, Mechanisms DB, and Flame DB. It uses EJB (Enterprise JavaBeans) to connect to these databases. For public databases provided by third-party providers, it uses http to request permission to access the databases (e.g. National Cancer Institute small molecule and PubChem databases). Thus the DB Client allows the user to access these databases and retrieve records to use in other tools, such as visualization for further modeling, without the necessity of manually downloading the records.

For private databases, CSE-Online currently has the Quantum Chemistry DB Management System (Quantum Chem DBMS) that allows a user to manage results from quantum chemistry calculations. Users can install (or have a System Administrator install) a Postgres or MySQL database system on their servers and create database accounts. The DB Client will allow the user to create private Quantum Chem databases on these servers for storage/query/retrieval of quantum chemistry results.

For shared databases, the DB Client allows users to share their private databases with other CSE-Online users (not necessarily having access to their servers) by copying these databases to the collaborative space on the CSE-Online server. The DB owner can set different levels of access, namely (Read only) or (Read and Write) for users with which they decide to share the databases.

How CSE-Online accesses computing resources

CSE-Online distinguishes three types of computing resources that the users have: those from their local desktop computers; remote servers where they have allocations; and from the computing grids such as TeraGrid.

Since CSE-Online is Java-based, it is running on the local user's CPU. In this case, graphic application tools in CSE-Online such as MolBuilder, BioViewer, or PsiViewer are using the local graphic cards on the user's desktop computers to achieve their performance. The virtual desktop continues to operate even in the event of Internet disruption. In fact, CSE-Online provides an offline mode where all local resources are accessible by Java applications in the CSE-Online environment.

For remote servers, both the GFE and the Resource Broker provide means for accessing files and application tools.

For a computing grid such as the TeraGrid, CSE-Online provides a Globus submission tool. This is a client interface using the Globus toolkit to submit jobs to any Globus enabled grid. Users must have a proper Certificate of Authority (CA) on the grid in order to use the Globus submission tool. It is different from many science gateways which have community accounts and submit jobs for users upon request from their portals. The Globus submission tool provides a direct means for users to submit jobs directly to the computing grid, not via any portal, and using their own CAs.

Work Anywhere At Anytime (WAAA)

As mentioned above, CSE-Online provides both online and offline modes so users can use the environment with their accessible resources at any time. In addition, CSE-Online allows users to export their current workspace configurations and preferences of all applications to an encrypted files which can be stored on a mobile storage

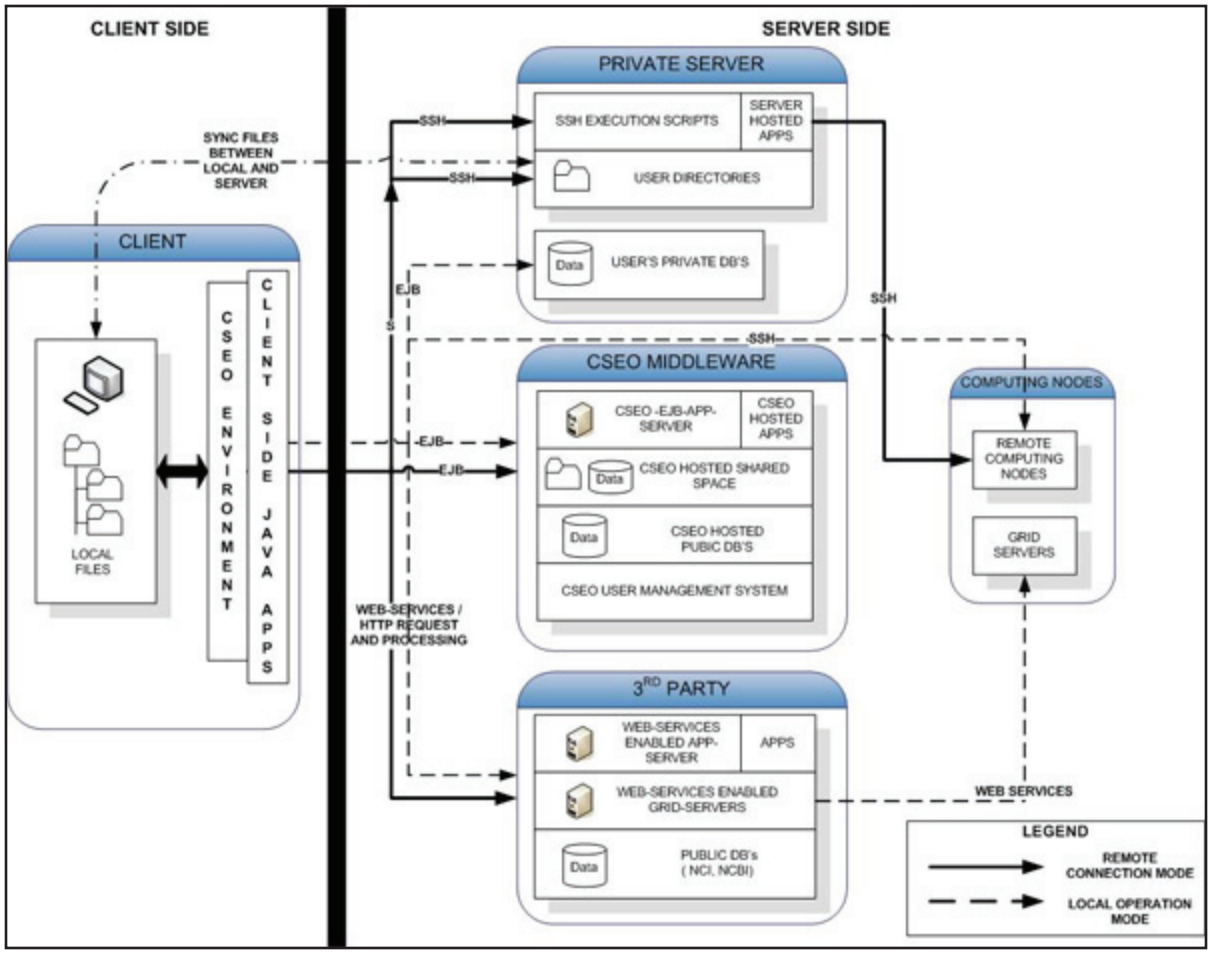

devices such as an USB drives. With this portable user workspace configuration, the user can activate the same working environment on any computer. This gives CSE-Online the WAAA characteristics.

The figure on the previous page illustrates the software architecture of CSE-Online; it provides a summary of the above discussion and shows how CSE-Online works.

CSE-Online currently has more than 30 applications in Quantum Chemistry, Chemical Kinetics, Bio-Simulation, and Material Science including a number of useful general purpose tools. The number of applications is growing steadily.

How to get started

Getting started on CSE-Online is rather easy with three simple steps:

1. Register for a free account at *http://cse-online.net*
2. Install CSE-Online on your desktop using the CSE-Online installer
3. Select applications from the catalog

CSE-Online.net also has a number of demos and tutorials. Please send comments on CSE-Online to: *truong@chemistry.utah.edu.*

CHPC Presentations

CHPC has developed a series of courses to help users make the most of their use of CHPC resources. We continuously add to and improve this series. We present these topics every Fall and repeat selected presentations again in the Spring. Please mark your calendars. These presentations are all held in the INSCC Auditorium and begin at 1:30 p.m. on the scheduled date.

RECENTLY PRESENTED:

*Overview of CHPC: 09/14/2006
*Intro Parallel Computing: 09/21/2006
*Intro to Programming with MPI: 09/28/2006
Intro to Programming with OpenMP: 10/12/2006
*Mathematical Libraries at CHPC: 10/19/2006
Debugging with Totalview: 10/26/2006

REMAINING TOPICS FOR FALL 2006:

11/2/2006: Parallel Performance Analysis with TAU

TAU (Tuning and Analysis Utilities) is a profiling and tracing toolkit for performance analysis of parallel programs. In this talk, we will introduce TAU as a new and flexible tool for tracing of parallel programs on CHPC Arches clusters. We detail small changes necessary to turn on the tracing and then explain how to visualize the trace files in Vampir trace viewer. We will conclude with some specific examples and glimpse on other features that TAU provides.

11/09/2006: *Chemistry Packages at CHPC

This talk will focus on the computational chemistry software packages - Gaussian, Amber, NWChem, Molpro, Gamess, Babel, GaussView, ECCE - that are available on CHPC computer systems. The talk will be an overview of the packages and their capabilities, and will focus on details of how users can access the installations at CHPC. This talk is the precursor for a second talk scheduled for next month that will focus on the use of Gaussian 98/03 and GaussView.

11/16/2006: *Using Gaussian03 and Gaussview

This presentation will focus on the use of Gaussian03 and Gaussview on the CHPC systems. The discussion will focus on the functionality of Gaussian as well as the format and construction of input files and PBS scripts. Restrictions on memory usage and disk space will be discussed. Timings of several jobs will be presented to demonstrate the parallel scaling that Gaussian achieves on icebox. Demonstrations on the use of GaussView will also be presented.

11/30/2006: Fast Parallel I/O at the CHPC

In this talk we explain how to perform fast parallel I/O operations on the CHPC computers. It should be beneficial for all users who are interested in speeding up their parallel applications via faster file operations. First, we describe in detail PVFS2 (Parallel Virtual File System) installed on Arches. Then we go over several examples on how to perform parallel I/O on this file system, in particular, MPI-I/O extension to the MPI standard and native PVFS function calls. Subsequently we detail ways to compile and run MPI-I/O applications on PVFS2. We conclude the talk with an insight into more advanced aspects of MPI-I/O.

Slides for presentations are available at: *http://www.chpc.utah.edu/docs/presentations/*.

**Presentations planned to repeat Spring 2007.*

FYI

CHPC maintains on its web site a listing of publications and talks that acknowledge the use of CHPC's resources. You can find the current listing at the following address:

http://www.chpc.utah.edu/docs/research/#chpcbib

If you utilize CHPC resources in your research, please include an acknowledgement in your publications and presentations (see page 8 for example). Also, please give us a copy for our records.

Article

iGrid Community Fileserver

by Julia Harrison and Brian Haymore

Center for High Performance Computing, University of Utah

CHPC recently put into production a community fileserver cluster. The idea is similar in concept to the "condominium" systems CHPC has implemented in the past with the icebox cluster and the current landscapearch cluster: CHPC will provide the infrastructure, and various research groups may purchase disk space, which may be used as departmental fileservers, large storage, and may also be mounted on the arches clusters.

The past fileserver support direction was not scaling well; it was becoming difficult to support further growth of so many independent fileservers, in both man-hours and time spent backing up these systems. It was functional and cost effective, but it was nearly impossible to keep a clean and concise plan to keep them in harmony with each other. Since they were not easily scalable, researchers found it simpler to buy additional servers rather than expand an existing fileserver. Department and users data sets were becoming more and more fragmented. When replacement servers were purchased, the downtime required to migrate data was excessive. Continuing in this direction could have meant dedicating several staff members just to do backups.

About a year ago we started exploring other options. The iGrid solution from Crosswalk Inc. (*www.crosswalkinc.com/*) presented itself and after testing this system for many months, we came to the conclusion that it was an excellent match. The design is modularized so, in a simple and clean way, you can increase desired services (i.e. memory, connectivity, and/or disk space). It also has many features which improve reliability, performance and maintenance.

Built into the system are several features we have found useful. The ability to perform data migration "hot" (while the system remains available) provides fewer interruptions of service to the user. In the existing model, if we have a server failure, the entire filesystem is out of service until we can make the necessary repairs. With the iGrid system two servers could fail and users would continue to have accessibility to all of the filesystems. Another improvement to maintenance is the ability to do rolling upgrades, where the servers are off-lined and upgraded one at a time, moving through the system so only one server at a time is unavailable, resulting in no interruption for the users.

Our current backup system performs a full backup on weekends and an incremental each weeknight. The full backups are in case of disaster recovery, where the incremental make it possible to recover files inadvertently removed or changed in the shorter term. The new iGrid system includes a built in snapshot feature, that could potentially eliminate the need for incremental backups. This could increase the amount of disk space we are able to backup, and/or increase the retention period of backups.

The iGrid system also has a built-in quota system.

For the past 8 months we have been using iGrid filesystems on our arches cluster to provide NFSroot services. We've experienced an increase in performance, and a significant reduction in the time it takes to boot the cluster (from 2 – 2 1/2 hours down to 15-20 minutes) both at planned and unplanned outages. The existing servers have usually 1 or 2 GigE connections where the new iGrid has 24 GigE connections which may be expanded.

The CHPC iGrid system hardware currently consists of 3 iGrid servers, each with 4 AMD Opteron processors and 16GB of RAM; 8 - 2 GBit/s fibre channel ports attached to the SAN; and 8 - 1 GBit/s Ethernet ports for network access via NFS/CIFS. Additional hardware includes 2 - 24 port McData 2 GBit/s fibre channel SAN switches; 2 Sun SE6130 controller trays; and 5 Sun SE6130 JBOD trays, for a total of 28 Terabytes of usable space (all allocated). As a beta tester CHPC has been pleased with the company's ability to address our concerns. The system comes with 24/7 support from the vendor. We have a minimum of 3 years of this support, which we may extend down the road.

We have also been able to maintain a nearly identical price point for researchers for purchasing standalone servers and space on this new cluster. This is possible because CHPC provided the current infrastructure. The system is expandable up to hundreds of terabytes.

If you are interested in more information about the iGrid system please contact us.

CHPC Staff Directory

Administrative Staff	Title	Phone	Email	Location
Julio Facelli	Director	556-2426	julio.facelli@utah.edu	410 INSCC
Julia D. Harrison	Associate Director	652-0019	julia.harrison@utah.edu	430 INSCC
Guy Adams	Assistant Director, Systems	554-0125	guy.adams@utah.edu	424 INSCC
Joe Breen	Assistant Director, Networking	550-9172	joe.breen@utah.edu	426 INSCC
DeeAnn Raynor	Administrative Officer	581-5253	dee.raynor@utah.edu	412 INSCC
Janet Ellingson	Administrative Assistant	585-3791	janet.ellingson@utah.edu	405-2 INSCC

Scientific Staff	Expertise	Phone	Email	Location
Victor Bazterra	Molecular Sciences	585-0003	victor.bazterra@utah.edu	405-32 INSCC
Thomas Cheatham III	Biomolecular Modeling	587-9652	tec3@utah.edu	306 INSCC
Martin Cuma	Scientific Applications	587-7770	m.cuma@utah.edu	418 INSCC
Byron L. Davis	Statistics	585-5604	byron.davis@utah.edu	416 INSCC
Julio Facelli	Molecular Sciences	556-2426	julio.facelli@utah.edu	410 INSCC
Stefano Foresti	Information Visualization	971-3680	stefano.foresti@utah.edu	405-13 INSCC
Robert McDermott	Visualization	581-4370	robert.mcdermott@utah.edu	420 INSCC
Anita Orendt	Molecular Sciences	231-2762	anita.orendt@utah.edu	422 INSCC
Ron Price	Grid Architect	560-2305	ronald.charles.price@gmail.com	405-4 INSCC

Systems/Network Staff	Title	Phone	Email	Location
Irvin Allen	System Administrator	231-3194	irvin.allen@utah.edu	405-40 INSCC
Thomas Ammon	Network Assistant	674-9273	tom.ammon@utah.edu	405-22 INSCC
Wayne Bradford	System Administrator	243-8655	wayne.bradford@utah.edu	405-41 INSCC
Erik Brown	System Administrator	824-4996	erik.brown@utah.edu	405-29 INSCC
Steve Harper	System Administrator	541-3514	s.harper@utah.edu	405-31 INSCC
Brian Haymore	Lead, Comp. Cluster Admin.	558-1150	brian.haymore@utah.edu	428 INSCC
Samuel T. Liston	Digital Communication & Visualization	232-6932	sam.liston@utah.edu	405-39 INSCC
Jimmy Miklavcic	Multimedia, Telematic & Digital Communication	585-9335	jimmy.miklavcic@utah.edu	296 INSCC
David Richardson	Computer Technician	550-3788	david.richardson@utah.edu	405-38 INSCC
Steve Smith	System Administrator	581-7552	steve.smith@utah.edu	405-25 INSCC
Matthew Thorley	Network Assistant	560-3438	matthew.thorley@utah.edu	405-20 INSCC
Neal Todd	System Administrator	201-1761	neal.todd@utah.edu	405-30 INSCC
Alan Wisniewski	Network Engineer	580-5835	alan.wisniewski@utah.edu	405-21 INSCC

User Services Staff	Title	Phone	Email	Location
Iris Price	Technical Assistant	N/A	iris.boanta@utah.edu	405-10 INSCC
Jason Duhaine	Systems Assistant	N/A	jason.duhaine@utah.edu	405-28 INSCC
Derick Huth	Technical Assistant	N/A	derick.huth@hotmail.com	405-19 INSCC
Beth Miklavcic	Multimedia Design, Digital Video	585-1067	beth.miklavcic@utah.edu	111 INSCC
Liza Newren	Technical Assistant	N/A	elizabeth.newren@yahoo.com	405-9 INSCC
Paul Vandersteen	Technical Assistant	N/A	paul.vandersteen@utah.edu	405-19 INSCC

The University of Utah seeks to provide equal access to its programs, services, and activities to people with disabilities. Reasonable prior notice is needed to arrange accommodations.

UNIVERSITY OF UTAH
Center for High Performance Computing
155 South 1452 East, RM #405
SALT LAKE CITY, UT 84112-0190

Welcome to CHPC News!

If you would like to be added to our mailing list, please fill out this form and return it to:

Janet Ellingson
UNIVERSITY OF UTAH
Center For High Performance Computing
155 S 1452 E ROOM 405
SALT LAKE CITY, UT 84112-0190
FAX: (801)585-5366

(room 405 of the INSCC Building)

Name:
Phone:

Department or Affiliation:
Email:

Address:
(UofU campus or U.S. Mail)

Thank you for using our Systems!

Please help us to continue to provide you with access to cutting edge equipment.

ACKNOWLEDGEMENTS

If you use CHPC computer time or staff resources, we request that you acknowledge this in technical reports, publications, and dissertations. Here is an example of what we ask you to include in your acknowledgements:

"*A grant of computer time from the Center for High Performance Computing is gratefully acknowledged.*"

If you use Arches, please add:

"partially supported by NIH-NCRR grant # 1S10RR17214."

Please submit copies of dissertations, reports, pre prints, and reprints in which the CHPC is acknowledged to: Center for High Performance Computing, 155 South 1452 East, Rm #405, University of Utah, Salt Lake City, Utah 84112-0190