

Case Research Computing Newsletter

23 February 2006

News from Case

1. Matching funds available from the Provost and VPITS/CIO for high performance computing
2. R is installed on the Case pilot HPC cluster
3. Mathematica is being tested on the Case pilot HPC cluster
4. MATLAB to be installed on the Case pilot HPC cluster
5. Software for computational chemistry to be installed on the Case pilot HPC cluster
6. Software for computational fluid dynamics to be installed on the Case pilot HPC cluster
7. Mass spectrometry/proteomics search engine to be installed on the Case pilot HPC cluster

Funding News

8. Daniel Atkins named head of new NSF Office of Cyberinfrastructure
9. Microsoft continues to fund Cornell Theory Center
10. Deadline to apply for NCSA summer faculty fellowships extended to March 15

News of Conferences and Professional Societies

11. O'Reilly Emerging Technology Conference - San Diego, California - March 6-9, 2006
12. Windows on the Future - Columbus, Ohio - March 7, 2006
13. Combined Sun HPC Workshop and Consortium Meeting - RWTH Aachen University, Germany - March 13-17
14. SAC 2006, Special Track on Distributed Systems and Grid Computing - Dijon, France - April 23-27, 2006
15. First International Conference on Grid and Pervasive Computing - Tunghai University, Taiwan - May 3-5, 2006 in conjunction with Grid Asia
16. Sun HPC Consortium Meeting - Singapore - May 14-15, 2006
17. Announcement and call for papers for the HPCS 2006 meeting.
18. Grid Asia 2006 IEEE International Symposium on Cluster Computing and the Grid - Singapore - May 16-19, 2006
19. JavaOne 2006 - San Francisco, California - May 16-19, 2006
20. IEEE International Symposium on High Performance Distributed Computing - June 19-23, 2006 - Paris, France
21. International Supercomputing Conference - Dresden, Germany - June 27-30, 2006
22. The Twelfth International Conference on Parallel and Distributed Systems - Minneapolis, Minnesota - July 12-15, 2006
23. Americas Conference on Information Systems (AMCIS 2006) - Acapulco, Mexico - August 4-6
24. International Conference on Parallel Processing - Columbus, Ohio - August 14-18, 2006
25. GridWorld - Boston, Massachusetts - October 3-6, 2006

Other news

26. National LambdaRail web site improvements

- 27. Seven things you should know about grid computing
- 28. IBM and Scripps Research Institute to collaborate on pandemic research
- 29. Global grid service for LHC computing succeeds in gigabyte per second transfer
- 30. TACC's top-ranked terascale compute cluster gets significant upgrade
- 31. Bioinformatics consortium adds supercomputing technology

1. Matching funds available from the Provost and VPITS/CIO for high performance computing

Matching funds continue to be available from the Office of the Office of the Provost and the Office of the VPITS / Chief Information Officer for researchers who have external funding for HPC equipment. Please see www.case.edu/its/researchcomputing/info/fund.htm for details. Note that this program has been extended since originally announced and that funds are still available.

2. R is installed on the Case pilot HPC cluster

The R statistical package (also known as GNU S) has been installed on the pilot HPC cluster. For more information on R, see www.r-project.org.

3. Mathematica is being tested on the Case pilot HPC cluster

The Mathematica software package has been installed and is in the process of being tested on the pilot HPC cluster. For more information on Mathematica, see www.wolfram.com.

4. MATLAB to be installed on the Case pilot HPC cluster

License acquisition is underway for MATLAB software for installation on the pilot HPC cluster. For more information on MATLAB, see www.mathworks.com.

5. Software for computational chemistry to be installed on the Case pilot HPC cluster

License acquisition is underway for two computational chemistry software packages for installation on the pilot HPC cluster: Gaussian and Amber. Other computational chemistry software is being investigated. For more information on Gaussian, see www.gaussian.com. For more information on Amber, see amber.scripps.edu.

6. Software for computational fluid dynamics to be installed on the Case pilot HPC cluster

License acquisition is underway for the FLUENT and FIDAP software packages for installation on the pilot HPC cluster. For more information on FLUENT and FIDAP, see www.fluent.com/software/fluent/index.htm and www.fluent.com/software/fidap/index.htm, respectively.

7. Mass spectrometry/proteomics search engine to be installed on the Case pilot HPC cluster

License acquisition is underway for the Mascot mass spectrometry / proteomics search engine software for installation on the pilot HPC cluster. For more information on Mascot, see www.matrixscience.com.

8. Daniel Atkins named head of new NSF Office of Cyberinfrastructure

<http://www.supercomputingonline.com/article.php?sid=10403>

University of Michigan professor Daniel Atkins has been appointed by the National Science Foundation (NSF) to head the newly created Office of Cyberinfrastructure. Atkins assumes his NSF role on June 1 and will oversee a \$182 million budget. The NSF created the Office of Cyberinfrastructure to make competitive, merit-reviewed awards for leading-edge, information technology-based infrastructure that is essential to science and engineering. Cyberinfrastructure includes supercomputers, data management systems, high capacity networks, digitally-enabled observatories and scientific instruments, and an interoperable suite of software and middleware services and tools for computation, visualization and collaboration.

9. Microsoft continues to fund Cornell Theory Center

<http://www.supercomputingonline.com/article.php?sid=10465>

Microsoft is funding a Microsoft Institute for High-Performance Computing at Cornell with annual funding of \$400,000, renewable each year for an indefinite period. The new institute will greatly expand the ability of researchers to work with huge databases of DNA sequences and protein composition and shapes, and explore new software and applications for the analysis of biological information. The institute will exist within the Computational Biology Service Unit (CBSU) at Cornell Theory Center (CTC). Its work will focus on the development of new software and applications in the Microsoft high-performance computing environment. Ron Elber, professor of computer science and director of CBSU, also will be director of the new institute, and it will be managed by CTC senior research associate Jaroslaw Pillardy, who

manages CBSU. Computational biology embraces the use of computers to simulate biological processes, such as the folding of proteins or the identification of gene function, as well as the management of large databases of biological information, an area known as bioinformatics. Biologists have created huge databases of DNA sequences and protein composition and shapes. To compare a sample with everything in these databases on an ordinary computer could take about a month, according to Pillardy, but a high-performance computer cluster like those operated by CTC can do it in minutes.

CBSU was created as part of the Tri-Institutional collaboration between biomedical and clinical researchers at Cornell, Rockefeller University and Memorial Sloan-Kettering Cancer Center in New York City, and has become a core resource for biology research throughout the Ithaca and New York campuses. It currently supports about 40 Cornell researchers. "We can give advanced computational services to biologists at significantly reduced cost, compared to what a single group can do," Pillardy said. "No single group would be able to afford these computational resources."

10. Deadline to apply for NCSA summer fellowships extended to March 15

<http://www.ncsa.uiuc.edu/summerff/>

The National Center for Supercomputing Applications (NCSA) will award six summer fellowships to faculty interested in contributing to the creation of a national cyberinfrastructure for science and engineering. The fellowships are available to researchers from across the country.

Faculty of the University of Illinois at Urbana-Champaign are not eligible (but can apply for the [UIUC/NCSA Faculty Fellows Program](#)). The deadline to apply has been extended to March 15.

The fellowships will provide 10 weeks' salary and will cover local housing expenses and the cost of travel to and from Champaign, Illinois.

Areas of interest include:

-- Cyberenvironments: NCSA is creating cyberenvironments that will simplify the integration of distributed computing and data resources into scientific and engineering research. We look forward to working with faculty who specialize in scientific and engineering applications, graphical user interfaces and portals, workflow and collaboration software, and data analysis and visualization tools.

-- Cyber-resources: NCSA provides high-end computing resources needed by the scientific and engineering communities. We want to spend the summer with faculty who are researching ways to make these resources better suit the communities they serve and to improve the computing infrastructure that they operate in.

-- Innovative Systems: Petascale computing is now a realizable goal that will impact all scientific and engineering research, but the best pathway to petascale computing is unclear. Come and explore those possibilities with us.

-- Cyberenvironments in the Classroom: Few of the advantages of a national cyberinfrastructure will be realized without scientists and

engineers who understand the new capabilities that it provides. NCSA is interested in collaborating with faculty to bring cyberenvironments into the classroom.

Applications are due March 15. For complete information and to apply online, go to <http://www.ncsa.uiuc.edu/summerff/>. Direct questions to ncsasff@ncsa.uiuc.edu.

11. O'Reilly Emerging Technology Conference - San Diego, California - March 6-9, 2006

<http://conferences.oreillynet.com/etech/>

Alpha geeks--hackers and other lead users--are a great early warning system for anyone who's watching the future of technology. The O'Reilly Emerging Technology Conference frames the ideas, projects, and technologies that the alpha geeks are thinking about, hacking on, and inventing right now into a coherent picture from which to extrapolate and upon which to start building.

12. Windows on the Future - Columbus, Ohio - March 7, 2006

<http://www.osc.edu/research/wof/>

Windows on the Future provides a forum in which academic institutions, research centers, corporate laboratories and government agencies can share knowledge, exchange ideas, and form networking relationships that will collectively benefit Ohio's higher education and industrial communities. Windows on the Future is sponsored by the [Ohio Supercomputer Center \(OSC\)](#), in cooperation with the [Ohio Board of Regents](#). This year there will be two tracks: Applying Advanced Networking and High Performance Computing to Health Information and Medical Research and Applying Broadband for Innovation in Education, Government, and Public Services. For more information, see <http://www.osc.edu/research/wof/>.

13. Combined Sun HPC Workshop and Consortium Meeting - RWTH Aachen University, Germany - March 13-17

<http://www.rz.rwth-aachen.de/sunhpc>

Combined Sun High Performance Computing Workshop and Consortium Meeting - SunHPC 2006, , RWTH Aachen University, Center for Computing and Communication, Aachen, Germany. Organized by the Sun HPC Consortium and the RWTH Aachen University

14. SAC 2006, Special Track on Distributed Systems and Grid Computing - Dijon, France - April 23-27, 2006

<http://www.cslab.ece.ntua.gr/sac06-dsgc/index.php>

This track aims to be a forum for scientists, engineers, and practitioners to share technical ideas, experience and results, and to present their latest findings in any aspects of distributed and grid computing.

15. First International Conference on Grid and Pervasive Computing - Tunghai University, Taiwan - May 3-5, 2006 in conjunction with Grid Asia

<http://hpc.csie.thu.edu.tw/gpc2006/>

Grid and Pervasive Computing (GPC) is an annual international conference on the emerging areas of merging grid computing and pervasive computing, aimed at providing an exciting platform and paradigm for all the time, everywhere services.

16. Sun HPC Consortium Meeting - Singapore - May 14-15, 2006

<http://www.sun.com/products-n-solutions/edu/he/researchandcomputation.html>

17. Announcement and call for papers for the HPCS 2006 meeting

<http://www.ace-net.ca/events/hpcs2006/>

The 20th International Symposium on High Performance Computing Systems and Applications will be held May 14-17, 2006 at Memorial University of Newfoundland in conjunction with OSCAR'06. See <http://www.ace-net.ca/events/hpcs2006> for more information. See <http://www.csm.ornl.gov/oscar06/> for information on OSCAR'06.

18. Grid Asia 2006 IEEE International Symposium on Cluster Computing and the Grid - Singapore - May 16-19, 2006

<http://pdcc.ntu.edu.sg/ccgrid2006/>

<http://www.ngp.org.sg/gridasia/2006/>

CCGrid 2006 provides researchers and practitioners with an opportunity to share their research and experience at the cross-roads of Grid Technology.

19. JavaOne 2006 - San Francisco, California - May 16-19, 2006

<http://java.sun.com/javaone/sf/>

The premier conference for Java developers.

20. IEEE International Symposium on High Performance Distributed Computing - June 19-23, 2006 - Paris, France

www.hpdc.org
<http://hpdc.lri.fr>

The Fifteenth IEEE International Symposium on High-Performance Distributed Computing (HPDC) will be a forum for presenting the latest research findings on the design and use of parallel and distributed systems for high end computing, collaboration, data analysis, and other innovative applications. Submissions are welcomed covering all aspects of high- performance distributed computing, Grids, and global computing ensembles. New scholarly research emphasizing empirical and reproducible results as well as investigative expositions of successful application and deployment efforts are particularly encouraged.

21. International Supercomputing Conference - Dresden, Germany - June 27-30, 2006

<http://www.supercomp.de>

The ISC is a premier venue for gaining an international perspective on the field of HPC. Combining a strong lineup of technical experts with exhibits from vendors, ISC presents state-of-the-art applications, architectures and trends in supercomputing.

22. The Twelfth International Conference on Parallel and Distributed Systems - Minneapolis, Minnesota - July 12-15, 2006

<http://www.icpads.umn.edu/>

ICPADS 2006 provides an international forum to exchange and share experiences, new ideas, and latest research results on all aspects of parallel and distributed systems.

23. Americas Conference on Information Systems (AMCIS 2006) - Acapulco, Mexico - August 4-6

<http://amcis2006.aisnet.org/>

24. International Conference on Parallel Processing - Columbus, Ohio - August 14-18, 2006

<http://www.cse.ohio-state.edu/~icpp2006/>

The conference provides a forum to present the latest research findings in any aspects of parallel and distributed computing.

25. GridWorld - Boston, Massachusetts - October 3-6, 2006

<http://www.gridworldhome.com/live/42/>

GridWorld is designed for business and technology professionals responsible for shaping directions and deploying grid solutions within research, industry and government enterprises.

26. National LambdaRail web site improvements

Bonnie Hurst, Project Manager at MCNC and NLR, announces that two sections have been added to the NLR website:

Services

<http://www.nlr.net/services/index.html>

The Services section provides an up-to-date description of NLR's current service offerings.

Researchers

<http://www.nlr.net/researchers/index.html>

The Researcher section is aimed at providing the research community with enough information to be able to propose utilizing NLR as the infrastructure to support their research projects. This online guide points them back to the "core set of services" and provides them with links to a list of each member's participants, administrative, and technical contacts and other relevant resources. Please review your member information for accuracy and let us know if you identify any discrepancies.

In addition, she points out a joint letter regarding NeTS support at <http://www.nlr.net/pubs/NLR-Internet2-NeTS2006.pdf>)

27. Seven things you should know about grid computing

<http://www.educause.edu/ir/library/pdf/ELI7010.pdf>

Seven things you should know about grid computing - from EDUCAUSE.

28. IBM and Scripps Research Institute to collaborate on pandemic research

<http://www.supercomputingonline.com/article.php?sid=10477>

On February 16, IBM and The Scripps Research Institute announced a collaborative initiative to conduct advanced research on pandemic viruses leveraging the industry-leading talent and technology from both organizations. The objective of "Project Check-mate" is to develop means to anticipate, manage and contain infectious diseases. The Check-mate Initiative will capitalize on Scripps Research's world class research in the areas of biochemical modeling and drug discovery and IBM's expertise in computational biology bio-patterning and supercomputing. The joint research team will exploit the use of IBM's Blue Gene supercomputer, the world's fastest supercomputer technology, and have access to Scripps Research's new state-of-the-art biomedical research facility, Scripps Florida, in Palm Beach County.

29. Global grid service for LHC computing succeeds in gigabyte per second transfer

<http://www.supercomputingonline.com/article.php?sid=10452>

On February 15, at the international Computing for High Energy and Nuclear Physics 2006 conference (CHEP'06) in Mumbai, India, the Worldwide LHC Computing Grid collaboration (WLCG) officially announced the successful completion of a service challenge involving sustaining a continuous flow of physics data on a worldwide Grid infrastructure at up to 1 gigabyte per second. The maximum sustained data rates achieved correspond to transferring a DVD worth of scientific data from CERN every five seconds.

The data was transferred from CERN in Geneva, Switzerland, to 12 major computer centers around the globe. Over 20 other computing facilities were also involved in successful tests of a global Grid service for real-time storage, distribution and analysis of this data. The completion of this service challenge is a key milestone on the way to establishing the necessary computing infrastructure for the Large Hadron Collider (LHC), the world's largest scientific instrument, which is scheduled to startup in 2007 at CERN. The results represent a significant step forward compared to a previous service challenge in early 2005 that had involved just seven centers in Europe and the USA, and achieved sustained rates of 600 megabytes per second.

The goal of the WLCG is to unite the efforts of established scientific Grid infrastructures to provide sufficient computational, storage and network resources to fully exploit the scientific potential of the four major LHC experiments: ALICE, ATLAS, CMS and LHCb. These experiments will be studying the fundamental properties of subatomic particles and forces, providing insight into the origins of the Universe. They are expected to generate in total some 15 million gigabytes of data each year. WLCG uses a range of national and international Grid infrastructures, including the Enabling Grids for E-Science (EGEE) project and the Open Science Grid (OSG).

30. TACC's top-ranked terascale compute cluster gets significant upgrade

<http://www.supercomputingonline.com/article.php?sid=10429>

On February 13, the Texas Advanced Computing Center (TACC) at The University of Texas at Austin announced a significant upgrade to its top-ranked terascale cluster, Lonestar. The upgrade, which will occur in two phases, will benefit researchers who rely on this powerful system to further research innovation in the areas of computational science, engineering and technology. Lonestar is the most powerful academic supercomputing system in the state of Texas and serves as a key resource in the [National Science Foundation TeraGrid initiative](#), the world's largest, most comprehensive distributed cyberinfrastructure for open scientific research. Cyberinfrastructure includes supercomputers, data management systems, high capacity networks, digitally-enabled observatories and scientific instruments, and an interoperable suite of software and middleware services and tools for computation, visualization, and collaboration.

31. Bioinformatics consortium adds supercomputing technology

<http://www.supercomputingonline.com/article.php?sid=10369>

To add significant computing and storage capacity to existing life sciences research infrastructure, the Bioinformatics Consortium at the University of Missouri recently purchased high-performance computing technology from Silicon Graphics and an SGI InfiniteStorage storage area network (SAN) with 8TB of capacity. The new SGI Altix server was specifically chosen for its shared-memory architecture and Linux operating environment, which will allow researchers, particularly in areas of molecular dynamics, molecular modeling, protein structure prediction, and a number of allied activities, to run larger-scale problems than ever before.

Installed in November, the SGI systems were made available to university researchers beginning in mid-December. The Altix will immediately take over a large number of big computational chemistry jobs currently being run on a slower server, including structural studies of bipolymers, chemistry in interstellar space, and polar order in crystalline organic molecular materials.

Researchers studying environmental issues using satellite and GIS data, and other scientists studying electrical properties of solid-state devices, have already requested blocks of time on the SGI Altix system as well. There is also very early-stage research of electronic structure, crystal structures and chemical structures of various biological compounds with the idea of being able to use natural organic compounds to develop computational capabilities. The primary reason the University purchased the SGI Altix system is to run these kinds of problems, which require massive processing power and shared-memory architecture to get results in hours instead of days or weeks.

About the newsletter

The Case Research Computing Newsletter is distributed on the last Thursday of each month and is intended to provide information of interest to researchers in the areas of high performance computing, high bandwidth networking, grid computing, visualization, and other uses of computation in research. To manage your subscription to the newsletter (or to unsubscribe), please visit <https://lists.case.edu/> and select the "RC-newsletter" mailing list. For comments about the newsletter or more information about research computing at Case in general, please contact Roger.Bielefeld@case.edu or see www.case.edu/its/researchcomputing. Old issues of the newsletter are archived at www.case.edu/its/researchcomputing/newsletter.