NSC News
2009:3
National Supercomputer Centre at Linköpings universitet

NSC'09 | 20 YEARS IN HPC 1989–2009

NSC | SNIC | PRACE

For more information and to register for the event, please visit http://www.nsc.liu.se/nsc09

Europe takes the lead

Europe to lead in high-end computing infrastructure

Access to high performance supercomputers is mandatory if European scientists and engineers are to remain internationally competitive. PRACE, the Partnership for Advanced Computing in Europe is preparing a pan-European HPC service as Thomas Eickermann, project manager, describes.

Read further on page 3

"Given the huge investments in high-performance computing worldwide, in particular in the USA and Japan, it is obvious that a coordinated European approach is mandatory – no single European country can meet the challenge alone."

NSC gets more international to meet future needs

It is with great pleasure that we welcome two new co-workers to Sweden, Linköping and NSC – Chandan Basu and Gabriel Mateescu. They will work with code optimisation and development issues, partly in the EU project IS-ENES devoted to computational infrastructure for climate modelling, and partly with other projects at NSC. They will have close contacts with our users to solve and optimise computational problems. Please read more about them in the detailed presentations in this issue of NSC News.

Chandan and Gabriel represent a new era at NSC, where we now increasingly also will provide expertise in code optimisation, especially with respect to future massively parallel systems. This expertise was hard to find within Sweden, so NSC this time had to extend our recruitment basis to India and USA. NSC will now have a better basis for increased user interactions and with the new expertise we pave the road for future developments within E-science.

How to construct and utilise future computer systems is also one of the topics at the upcoming conference NSC'09 which takes place here in Linköping on 13 and 14 October. One part of the conference will focus on PRACE – the EU-project aiming at pre-studies for the next generation of large-scale computers. One part will be the annual SNIC Interaction, where users present computational solutions to various scientific problems. A third part of the conference will be devoted to flashbacks in high-performance computing in Sweden during the last 20 years. In conjunction with the conference there will also be a workshop on large-scale storage. More details will be posted on the NSC web site. We wish you all very welcome to the NSC'09 event.

Finally, as autumn arrives, we will initiate the test phase for our new capacity system, which we plan to have up and running by the end of the year. This new system will be devoted both to SNIC users from all parts of Sweden and to local users at Linköping University. We look very much forward to this next step in capacity increase at NSC.

BENGT PERSSON, NSC DIRECTOR

Kristin Winader makes it again

NSC continues to artistically decorate our computer hall. This time it is our main workhorse Neolith that has received the attention from Kristin Winander, who was inspired by the Neolith job scheduler status display. She is an artist that works with a wide variety of materials besides painting. And has had several exhibitions in Linköping, Norrköping and Stockholm.

New staff member: Chandan Basu

Hi, my name is Chandan Basu. I have done Ph.D. in Computational Physics from Indian Institute of technology, India. I have been working on parallel applications for several years. I have earlier worked on scaling / tuning of applications on large clusters, benchmarking of large clusters, network topologies of large clusters etc. I have joined NSC as an expert on distributed computing. I will be working on optimization, benchmarking and other applications and system related issues.

Europe to lead in high-end computin

Over the last three decades, computational science has evolved as the 'third pillar' of scientific research complementing theory and experiment. Scientists and engineers rely on computer simulations, when the problems are too complex to be solved by theory alone or when experiments are too dangerous, too expensive, or simply impossible. Some problems can be solved by distributing the work over many smaller systems.

However, each scientific discipline has tightly coupled problems that require the capabilities of a single high-end computing resource. In 2006, leading European scientists from climatology, weather and earth sciences, astro-, high-energy and plasmaphysics, materials science, chemistry and nanoscience, life sciences and engineering have produced a scientific case estimating that in 2009, European researchers need access to leading computing systems of 1 Petaflop/s peak performance to remain internationally competitive (1 Petaflop/s corresponds to 10^{15} Floating point operations per second). The advent of the first Petaflop/s computer in the summer of 2008 has proven them right.

Given the huge investments in high-performance computing worldwide, in particular in the USA and Japan, it is obvious that a coordinated European approach is mandatory – no single European country can meet the challenge alone. This has also been acknowledged by the European Strategy Forum on Research Infrastructures (ESFRI) that has identified High Performance Computing (HPC) as a strategic priority for Europe in its 2006 roadmap.

PRACE for world-class computing

PRACE, the Partnership for Advanced Computing in Europe, was formed in 2007 as an initiative of 14 – and subsequently 20 – European countries with the goal to implement the ESFRI vision of a world-leading, persistent, pan-European high-end computing infrastructure. This infrastructure is to be managed as a single European legal entity. It will comprise several world-class supercomputer centres offering a range of architectures to meet the needs of the different scientific and industrial domains and applications. In a 'preparatory phase' project, funded in part by the EC under FP7 grant agreement n° RI-211528, PRACE is undertaking all legal, administrative and technical work to establish the infrastructure to start its operation in 2010.

This two year project commenced in January 2008 with an overall budget of 20 million Euros, out of which 10 million Euros are contributions by the EC. Its tasks include the selection of an appropriate legal form, the definition of a robust and flexible governance structure that manages the relations between the national shareholders, the European Commission, the hosting centres and the scientific users. Ensuring sustained funding is a key objective. Access will be provided through a Europe-wide peer review process, following the principle 'world-class science on world-class systems'.

The European HPC ecosystem

Today, HPC services are provided typically on the national level. Most

New staff member: Gabriel Mateescu

I have joined NSC in September 2009 as a computational scientist for High Performance Parallel Computer Systems. My work at NSC and my background are in the development of infrastructure for accessing and managing computational and data resources used or produced by a wide range of applications, from bioinformatics to computational chemistry and climate modeling.

ng infrastructure

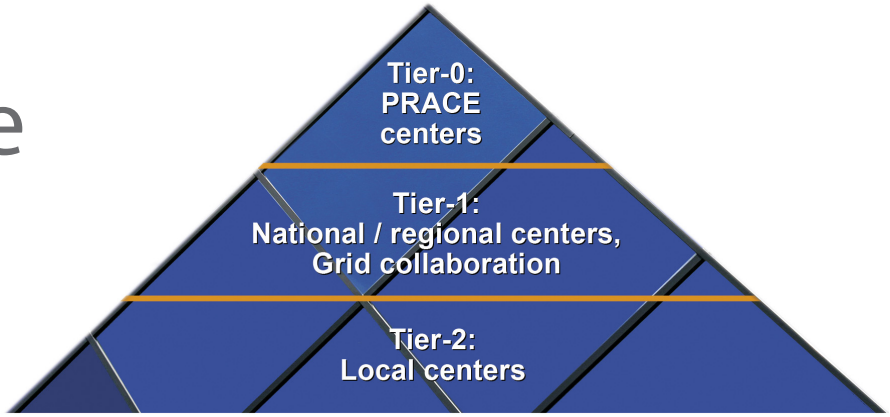

countries are operating national HPC centres (so-called tier-1 centres) to meet high-end demands that cannot be fulfilled by regional and local centres (tier-2). PRACE will add a European tier-0 to this HPC ecosystem that is often visualised by a performance pyramid (see Fig. 1). Here the vertical axis represents computational capability and the horizontal axis indicates the number of systems in the layer. Close coordination with the tier-1 centres will maximise the benefit of the new tier-0 to the users, allowing a smooth transition of users between national and European systems. In particular, PRACE is closely collaborating with DEISA, the Distributed European Infrastructure for Supercomputing Applications. DEISA is coordinating European tier-1 resources and offers a small part of their computing resources for European projects. PRACE and DEISA are working together on middleware developments for ecosystem integration. Both projects have close links with EGI_DS and EGEE with the goal to ensure Grid access and interoperability between the infrastructures.

Technical achievements

The legal and administrative work of the PRACE project is complemented by technical and scientific work to prepare the timely operation of the research infrastructure in 2010. The latter focuses on three areas: analysis of key applications and scaling of related codes; selection and evaluation of architectures for likely Petaflop/s production systems; identification of training needs and education in advanced programming methods.

To prepare the procurements of the future tier-0 machines PRACE conducted a survey of the utilisation of HPC resources across Europe. This resulted in a core list of applications which represent the major scientific disciplines, contain the core algorithms, and use most of the cycles of the surveyed systems. The derived benchmark suite will be used in the procurement of the tier-0 systems. In parallel, PRACE conducted an extensive technology survey, involving all major HPC vendors, to identify technologies that are likely to deliver Petaflop/s performance in 2009/2010.

Matching the application requirements to the technology options resulted in the selection of six prototype systems that are being used to evaluate the architectures and the scalability of the codes. This comprehensive set of architectures consists of two massively parallel systems (Cray XT5, IBM Blue Gene/P), a thin node cluster (Intel Nehalem processors + Infiniband), a fat node cluster (IBM Power6) and two hybrid architectures (IBM Cell + Power6, NEC SX-9 + Intel Cluster). More advanced technologies for multi-petascale systems beyond 2010 are investigated in close collaboration with vendors. This 'bootstraps' a continuous process of technology evaluation to foster HPC developments in Europe in the long run.

To effectively exploit future systems new programming paradigms are needed. PRACE conducted a survey of HPC training needs among the user communities. Based on the results a very successful summer school was

New computing power to NSC

NSC will install a new HPC cluster this autumn as a capacity complement to the capability system Neolith. Half of the capacity of the new system will be dedicated to users at Linköpings universitet while the other half will be distributed through SNAC. The system will be delivered by HP through GoVirtual and is planned to be shipped in mid October. It will consist of 362 compute nodes and 2 system nodes. The compute nodes are HP SL170h G6 servers with two Quad-core Intel Xeon 5520 processors. A majority of the compute nodes will be equipped with 24 GiB primary memory and 56 'fat-nodes' will have 72 GiB of primary memory. Two networks will be available on the system, a Gigabit Ethernet and a low capacity Infiniband (4x1 oversubscribed).

PETER MÜNGER

held in Stockholm in August 2008 followed by a winter school in Athens in February 2009. This is the start of a permanent training and education activity of PRACE.

The PRACE project partners from 20 European countries are convinced they will succeed in their challenging mission to establish an HPC service as a persistent European research infrastructure in just two years.

Contact Details

Forschungszentrum Juelich
Prof Achim Bachem
PRACE Coordinator
52425 Jülich, Germany
T: +3589 457 2058
E: prace-coordinator@fz-juelich.de
W: www.prace-project.eu

Dr. Thomas Eickermann is head of the Communication Systems division at the Jülich Supercomputing Centre at Forschungszentrum Jülich, Germany.

Do you know where your password is?

Lately, a large number of academic institutions world-wide have been the target of computer intrusions through the use of stolen identities. Each such incident incurs a large cost in lost productivity and clean-up efforts, both for the institutions themselves and for the individual users. Just imagine if you had to get new passwords for every single account you use, or if your research data was deleted or – perhaps worse – tampered with.

Given the current threat situation, it is more important than ever that you keep your login credentials – passwords, crypto keys, certificates – safe and secure.

The most common authentication system is of course ordinary passwords, and you probably already know how passwords should be handled. Still, let us underline a few important points:

- Choose good passwords. One method for this is to choose a sentence that is easy to remember, and then combine the initial letters in each word to a password, perhaps with a few twists. For example, the Shakespeare quote "Threescore and ten I can remember well" can be used to construct the password "3SatIcrw".
- Keep your passwords secret. By all means, write your passwords down on a (well-guarded) piece of paper, but never store them in clear-text on your computer, and never share them with anybody. Any NSC user found to have shared his password gets his account immediately terminated – no excuses, no exceptions.
- Use different passwords. If an intruder manages to steal one of your passwords, he shouldn't get access to all of your accounts – use different passwords for different systems. Again, it is better to have to write your passwords down than to use a single password for all systems.

There are other authentication systems than passwords, though. Many sites, including NSC, support and encourage the use of ssh keys to authenticate to their systems – please see NSC's ssh user guide at http://www.nsc.liu.se/support/userguides/remote/ssh.html for details.

Properly used, ssh keys are actually both more secure and more convenient than passwords. Unfortunately, they can also be very insecure, if used improperly.

The single most important thing about ssh keys is that you must protect your private ssh key with a good passphrase. Preferrably, your private ssh key should never leave your personal computer.

Fortunately, the use of the ssh-agent program together with the agent forwarding functionality of the ssh client make this easy to set up – this is also covered in NSC's ssh user guide.

The security of NSC's systems and the integrity of your data is in your own hands. Please take due caution.

LEIF NIXON, SECURITY OFFICER

UPCOMING EVENTS

ICPP-2009; 38th International Conference on Parallel Processssssing
September 22 – 25, 2009, Vienna, Austria.
http://www.cse.ohio-state.edu/~icpp2009

NSC'09 Conference with PRACE Code Porting Workshop and SNIC Interaction
October 13 – 14, 2009, Linköping, Sweden.

Topics:
- presentation of PRACE prototype systems
- examples of PRACE systems usage
- demonstrations of code porting issues
- user presentations
- 20 years of Swedish HPC

http://www.nsc.liu.se/nsc09

DS-RT2009; 13th IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications
October 25 – 28, 2009, Singapore.
http://www.cs.unibo.it/ds-rt2009/

HEPiX Workshop
October 26 – 30, 2009, Berkely, CA, USA.
http://indico.cern.ch/conferenceDisplay.py?confId=61917

LISA'09; 23rd Large Installation System Admininstration Conference
November 1 – 6, 2009, Baltimore, MD, USA.
http://www.usenix.org/event/lisa09

SC09; International Conference for High Performance Computing, Networking, Storage and Analysis
November 14 – 20, 2009, Portland, Oregon, USA.
http://sc09.supercomputing.org

HiPC 2009; 16th IEEE International Conference on High Performance Computing
December 16 – 19, 2009, Kochi, India.
http://www.hipc.org

FAST'10; 8th USENIX Conference on File and Storage Technologies
February 23 – 26, 2010. San Jose, CA, USA.
http://www.usenix.org/events/fast10

Per-Com 2010; 8th IEEE International Conference on Pervasive Computing and Communications
March 29 – April 2, 2010, Mannheim, Germany.
http://www.percom.org

IPDPS 2010; 24th IEEE International Parallel & Distributed Processing Symposium
April 19 – 23, 2010. Atlanta, Georgia, USA.
http://www.ipdps.org

CCGrid10; 10th IEEE International Symposium on Cluster Computing and the Grid
May 2010, Melbourne, Australia.
http://www.gridbus.org/~raj/ccgrid

Linköpings universitet

National Supercomputer Centre, Linköpings universitet, 581 83 Linköping, Sweden
Tel: +46 (0) 13-28 26 18, fax: +46 (0) 13-28 25 35, e-mail: info@nsc.liu.se
www.nsc.liu.se

Editor: Peter Münger To subscribe, for a paper version and/or PDF-file, send an email to info@nsc.liu.se