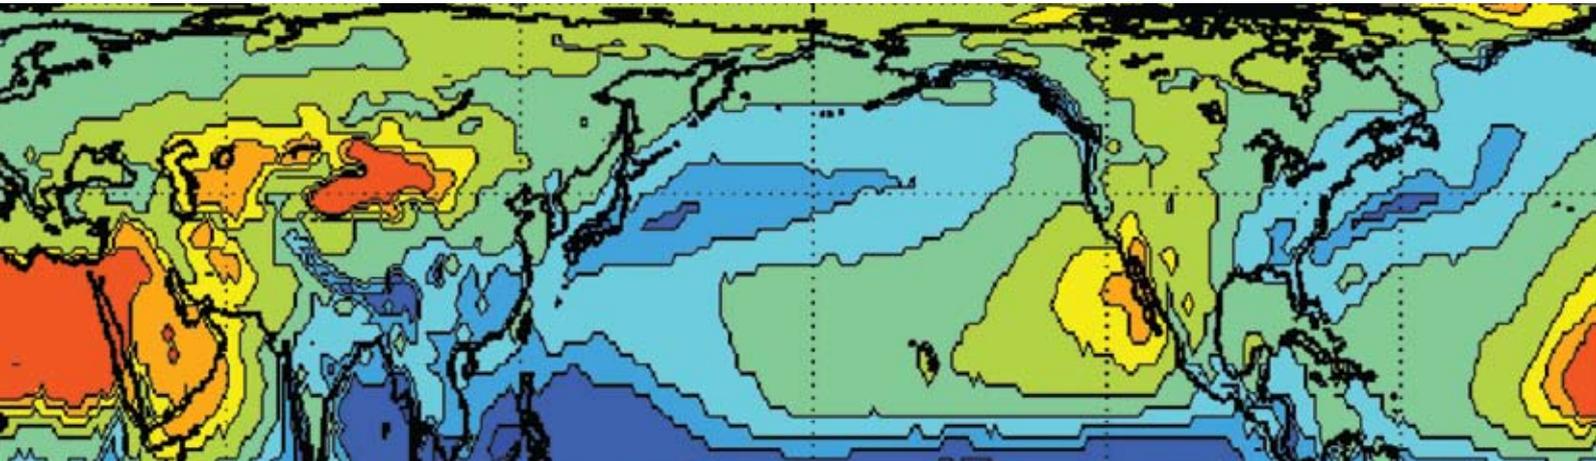




KTH Computer Science
and Communication

no:1
2009

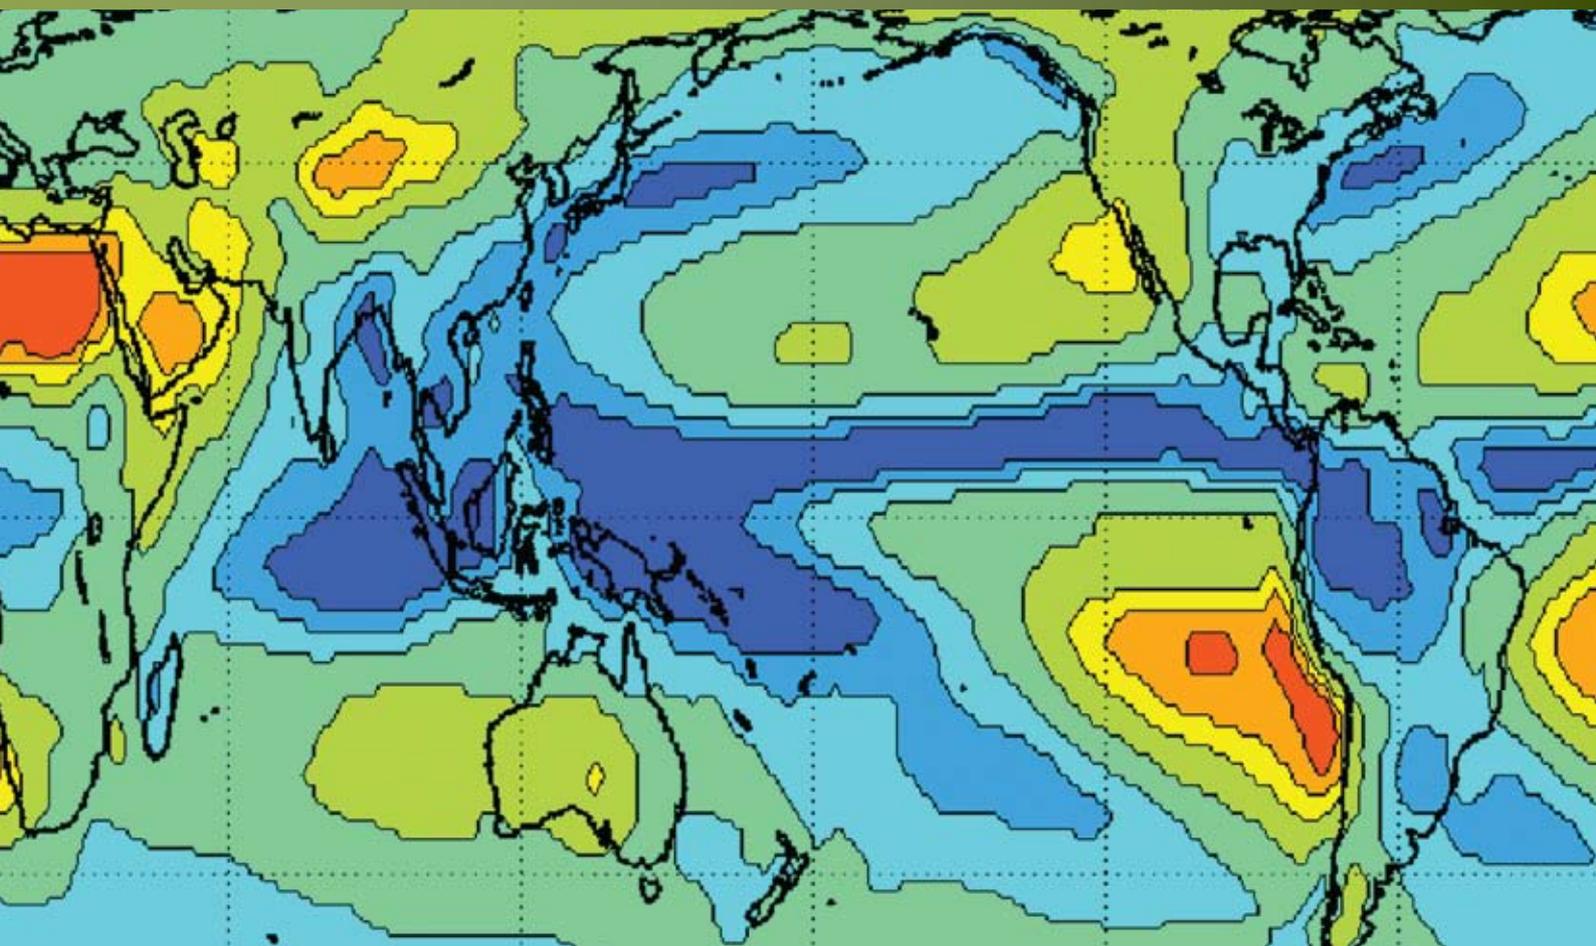


PDC Newsletter

Plans for Coupled Global Climate Modelling on Ekman and Vagn:

EC-EARTH by Klaus Wyser and Colin Jones, Rossby Centre, SMHI

Can We Beat Batman? by Elisabet Molin, PDC





LENNART JOHNSON,
DIRECTOR PDC



ERWIN LAURE,
DIRECTOR PDC-HPC

PUBLISHED BY PDC AT KTH CSC

The center operates leading-edge, high-performance computers as easily accessible national resources. These resources are primarily available for Swedish academic research and education. PDC receives funding from the Swedish Research Council (VR), and KTH.

Editors:

Lennart Johnsson and Erwin Laure

Editorial staff: Mike Hammill

Layout & Graphic Design:
Maria Malmqvist

E-mail: pdc-newsletter@pdc.kth.se

ISSN 1401-9671

With this issue of the PDC newsletter we are reviving the PDC newsletter tradition that had become a bit dormant. Our goal with the newsletter is to inform you of the latest developments at PDC in infrastructure, systems and their operations, and research and development. These developments are aimed at improving the environment and services we provide as well as cost effectiveness, training and education opportunities and offerings. An equally important goal is to give you insights into the exciting science that is being pursued using PDC managed resources.

PDC has seen many changes during the past year in terms of personnel, infrastructure and systems. PDC's deputy director, Per Öster, moved to the Finnish IT Center for Science, CSC, Helsinki, to assume the role as Director for Application Services. After an interim period where Peter Graham, the former head of KTH's Network Operations Center, KTHNOC, served as deputy director, KTH was fortunate to have Erwin Laure join PDC as director of PDC's HPC infrastructure and operations activities, called PDC-HPC. Erwin joined KTH from CERN, the European Organization for Nuclear Research in Geneva, Switzerland, where he served as Technical Director of the EGEE (European Grid for E-Science) EU-funded project. In this position, Erwin was responsible for coordinating the operations, middleware, and user support of the largest multidisciplinary Grid infrastructure in the world. Earlier, some long-term systems and user-support staff members, familiar to many of PDC's users, have left for positions in industry. Recruiting to fill the vacancies is still in progress, though some positions have been filled. These changes have taken place at a time of exceptional growth in the compute and storage resources managed by PDC, and the required supporting infrastructure.

The exceptional growth in capabilities is well illustrated by the total capability of systems managed by PDC: in Fall 2000, it was 0.217 teraflops (TF); in Fall 2008 it was 173 TF. This represents an 800-fold increase, well above the increase in capability predicted by Moore's law (45-fold over the same period). However, the most dramatic increase took place during the approximately 15 months from Summer 2007 to Fall 2008, a period in which the capability of PDC systems increased from about 16TF to the current 173TF. During this period a 5,440 core, 57TF Intel Harpertown system called Ferlin, a 10,144 core, 90TF AMD Barcelona system called Ekman, and a 744 core, 7.7 TF Intel Harpertown system for SweGrid became available. Ferlin is in part a SNIC Foundation Level System intended for high-throughput applications that do not require fast interconnect, with 237 of the nodes (1896 cores) owned by the Stockholm Center for Biomembrane Research. Ekman, currently the most powerful public supercomputer in Sweden, is jointly owned by Stockholm University, KTH, and SMHI and is dedicated for climate research and flow simulations. Ekman is operated in collaboration with NSC in

Linköping who host a storage system for climate research, called Vagn, used both for Ekman and SMHI computer systems at NSC and SMHI. Last, but not least, we have also upgraded our tape library and installed a new IBM TS3500 tape robot with a current capacity of 2.3 Petabytes.

The more than 10-fold increase in compute power since early 2007 to the middle of 2008 required about a three-fold increase in PDC's electricity, UPS and cooling infrastructure that now is at a level of 1MW. This increase represents an investment by KTH at a level of close to 20 Mkr with significant annual costs for energy consumption and heat removal (estimated at more than 8Mkr with current prices), a clear sign of the importance KTH attributes to computational science and engineering for its mission and vision. In the near future, PDC's network capability will also be significantly enhanced.

Beginning with this newsletter we start a series of articles that highlight the science that is being enabled by the resources at PDC. The first such article is on EC-EARTH, a new coupled atmosphere-ocean model that will run on Ekman. Also in this newsletter we start a series on user support at PDC, and introduce new staff members. We also plan to report on user interaction events, such as education and training opportunities, and the "PDC Pub" that was held on April 6th for the first time. This event was hosted by PDC together with the KTH Computational Science and Engineering (KCSE) Graduate School to facilitate communication between system users and our staff. We plan to make this a recurring event.

Lennart Johnsson, Director PDC

Erwin Laure, Director PDC-HPC

Photo: Harald Barth



PDC-KCSE PUB, APRIL 6

The PDC-KCSE Pub is a great way to meet PDC staff in an informal setting. Pictured above, participants of the Pub assemble in PDC's kitchen to chat and ask questions of PDC's staff while having a free beer and sandwich. PDC's Elisabet Molin, foreground, hands out a brief survey. Participants were also invited to visit PDC's computer hall.

In This Issue:

PDC EDITORIAL

page 2

PLANS FOR COUPLED GLOBAL CLIMATE MODELLING ON EKMAN AND VAGN: EC-EARTH BY KLAUS WYSER AND COLIN JONES, ROSSBY CENTRE, SMHI

page 3

CAN WE BEAT BATMAN? BY ELISABET MOLIN, PDC

page 6

STAFF FOCUS

page 5,6,7

PDC RELATED EVENTS

page 8

SOURCES

page 8

On the Cover:

*"Annual Mean Accumulated
Precipitation as simulated by EC-EARTH"*

HOW TO CONTACT PDC

Visiting address:

Teknikringen 14, plan 5, KTH, Stockholm

Mail: PDC, KTH,
SE-10044, Stockholm, Sweden

E-mail: support@pdc.kth.se

www: <http://www.pdc.kth.se>

Phone: +4687907800

Fax: +468247784

PDC Related Events

Introduction to High-Performance Computing

Photo: Mike Hammill



This year's PDC Summer School takes place August 17-28 on the KTH main campus.

This is an intensive two-week course for a limited number of participants.

A number of topics will be covered in overview talks, in-depth technical lectures, and hands-on computer lab sessions.

For more information, visit the Web site: <http://www.pdc.kth.se/education/summer-school>

Learn about Grid and hear PDC's Åke Edlund give a lecture on clouds computing at ISSGC09.

ISSGC09 – INTERNATIONAL SUMMER SCHOOL ON GRID COMPUTING 2009

July 5-17, 2009, Nice, France

<http://www.iceage-eu.org/issgc09/index.cfm>

Plans for coupled global climate modelling on Ekman and Vagn: EC-Earth

By Klaus Wyser and Colin Jones, Rossby Centre, SMHI

EC-EARTH is a new, coupled, atmosphere-ocean model that consists of models for the atmosphere (known as Integrated Forecast System IFS) and ocean (NEMO) that communicate through a coupler (OASIS3). The model is based on the seasonal prediction system of the European Center of Medium-Range Weather Forecasts, and adapted to the specific needs for climate simulations. EC-EARTH is jointly developed by a consortium of European MetServices and Universities. Currently, final tests are under way to check the model's ability to simulate a realistic present-day climate.

The coupled model intercomparison project CMIP5 (<http://www-pcmdi.llnl.gov/projects/cmip/index.php>) is a new attempt to coordinate a series of experiments with state-of-the-art climate models. Climate models that are willing to contribute to CMIP5 are required to follow a pre-defined protocol that specifies the details of a climate simulation such as the initial conditions or the change of the external forcing (e.g., greenhouse gases) with time. Results from the previous CMIP3 have been used for the United Nations Intergovernmental Panel on Climate Change (IPCC) 4th assessment report. Likewise, CMIP5 is intended to produce a huge database supporting the IPCC work for the next assessment report (AR5) that is planned for 2014.

Two streams are planned for CMIP5. In the classic climate experiment stream, the models start from pre-industrial conditions and simulate the evolution of the climate until today following observed external forcing. These runs will then be extended into the future with forcings from specific scenarios that describe potential pathways for future development and associated greenhouse-gas concentrations. The second stream of experiments is novel and geared towards climate predictions for the next 10-30 years. The biggest challenge with climate predictions is the creation of initial conditions for the ocean, which is the part of the climate system that evolves relatively slowly and might retain a predictable signal on decadal timescales. Within CMIP5, a large number of decadal hindcasts will be made, whereby a climate model is initialized from a certain date and then run over a number of years within the recent past. In doing this it is then possible to use observations to assess the quality of the model simulations and gauge the degree of predictability that remains as the integration proceeds. Climate prediction is one of the leading edges of climate research today, with many outstanding issues remaining to be resolved.

The Rossby Centre at SMHI plan to use EC-EARTH for both streams of CMIP5. This requires a strong computing infrastructure as each CMIP5 experiment consists of 100 or more years of climate simulations with complex, coupled climate models. The output of each experiment is on the order

of several TB that need to be post-processed and analyzed. Fast computers and high bandwidth are therefore crucial for the planned EC-EARTH contribution to CMIP5. The new Ekman cluster and Vagn storage facility are welcome additions to the Rossby Centre's computing environment and will enable the Rossby Centre to make a significant contribution to the CMIP5 project.

As an initial example of some results from EC-Earth integrations Figure 1 shows the annually averaged accumulated precipitation as simulated by EC-Earth when observed Sea Surface Temperatures (SSTs) and Sea-Ice fractions are prescribed as a lower-boundary condition for the model. Results are presented as a climatological average for the period 1990-1999, inclusive- Compared to the satellite observed values derived from the GPCP product (Adler et al. 2003), EC-Earth successfully simulates the majority of the main areas of intense rainfall. In particular, extreme maxima along the equatorial Pacific and into the West Pacific and Indian Ocean are well captured, the so-

...continued on page 6

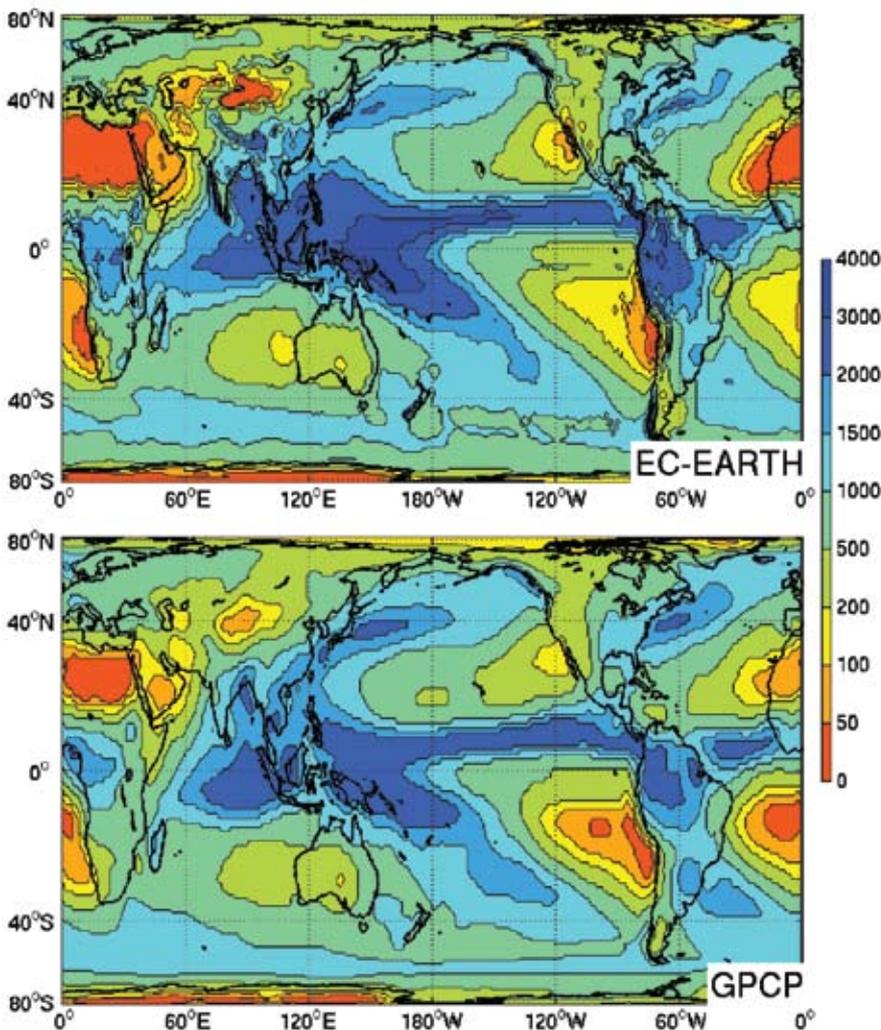


Figure 1. Annual Mean Accumulated Precipitation (mm) as simulated by EC-Earth and from the GPCP (Global Precipitation Climatology Project) for the period 1990-1999 inclusive.

Staff Focus



MICHAELA LECHNER

Michaela Lechner has studied Physics and Astronomy at the Institute of Astro and Particle Physics in Innsbruck, Austria, and was employed several years within the Austrian Grid project [<http://www.austriangrid.at/>]. She started at PDC in autumn as a part-time replacement for application expert Ulf Andersson. After finishing her PhD in November she began working full time at PDC on January 1st. During her first months she mainly contributed to maintenance and setup, especially of the next SweGrid cluster, and honed her Grid-operation skills. Her new focus is to provide better overall application support by working more closely with users. In particular, she will be the contact person for the astrophysics community. Her goal is to facilitate and optimize the usage of PDC's HPC systems and the Grid for the users. If you have further questions don't hesitate to write an email to <caela@pdc.kth.se> directly.

Says PDC director, Erwin Laure: "Michaela will be of great help to our users. Astrophysicists in particular will find that with her help on code speedup, 'the sky is the limit!'"

Staff Focus



DEJAN VITLACIL

Dejan Vitlacil is from Sarajevo, Bosnia and Herzegovina. He graduated with an MSc in Corso di laurea in Scienze dell'Informazione (computer science) [<http://www.eng.unibo.it/PortaleEn/Academic+programmes/Courses/MatematicalSciences/2ndDegree/2005/CoursePage20050246.htm>] at Universita' degli Studi di Bologna in Italy in 2003. While working as a Software Developer on an eGovernment project in Italy, he graduated again with an MSc in Corso di Laurea Specialistica in Informatica at Universita' degli Studi di Bologna in 2006. After a year's experience as a system administrator in a Windows environment, he began working at the Italian National Center for Research and Development about Information and Data transmission Technologies (CNAF) [<http://www.cnaf.infn.it/main/index.php/en/>]. CNAF is the National Center of the National Institute for Nuclear Physics (INFN) [<http://www.infn.it/indexen.php>]. He did R&D in the field of information technologies as a research fellow, with the position of Storage Group Administrator. His work was mainly centred on the management, configuration, maintenance and development of TIER1 Storage apparatus. In particular, he worked on CERN Advanced STORage manager (Castor) [<http://castor.web>

...from page 5

called Inter Tropical Convergence Zone (ITCZ) and its southward extension into the southwest Pacific, the South Pacific Convergence Zone (SPCZ). Likewise, maxima over central Africa and South America are also well simulated. Increased rainfall rates can be observed in the mid-latitude storm tracks over both the Atlantic and Pacific oceans. The Pacific Ocean storm track precipitation appears somewhat weak in terms of absolute intensity and may be linked to an underestimate in intensity and/or numbers of mid-latitude cyclones which are the main rain-bearing systems in this region.

Relative minima in precipitation are correctly simulated, both with respect to location and intensity, over the major arid regions of the subtropical western continents and the eastern subtropical oceans. Discrepancies in relation to the rainfall intensity along the Himalayan foothills, with significantly more precipitation in the EC-Earth simulation than in the GPCP observations, may partially be a result of the low resolution of the satellite data set ($2.5^{\circ} \times 2.5^{\circ}$), being unable to resolve the intense rainfall in this region which is supported by other, higher-resolution, land-based observational data sets. The main area of disagreement appears to be in the location and intensity of the Southern Hemisphere circumpolar storm track, the relative maximum in precipitation seen around the Southern Ocean at $\sim 50\text{-}55^{\circ}\text{S}$. This feature appears somewhat weaker, less continuous and offset to the north by $\sim 5^{\circ}$ in the EC-Earth simulation. Rainfall in this region is associated with mid-latitude cyclones, one possible explanation for the underestimate, analogous to the problems in the North Pacific, is therefore an underestimate of the number and/or intensity of these systems in the EC-Earth simulation. This issue will be analyzed in the near future.

It is important to remember that the results presented here are with prescribed SSTs and sea-ice, runs made for the full CMIP5 experiment require that the model is run in coupled mode and actually simulates the evolution of the ocean and resulting SSTs and sea-ice amounts. This is a far more challenging task and one can necessarily expect that simulated precipitation rates will be inferior to those presented here, in particular where errors in SSTs strongly force precipitation bearing circulations. Results from the coupled CMIP5 runs will be carefully analysed in the coming year at the Rossby Centre.

References:

Adler R.F., Huffman G.J., Chang A et al. 2003. The Version-2 Global Precipitation Climatology Project (GPCP) monthly precipitation analysis (1979-present). J. Hydrometeor, 4 (6), 1147-1167

Can We Beat Batman?

by Elisabet Molin, PDC

Gotham City is in danger. The chief of police flashes the bat light on the sky--help! What is the emergency? Is it a new villain? Or a trouble maker revisited, the Joker? Or the Riddler maybe?!



Sketch picture:
Elisabet Molin

At PDC we have a system that makes it possible for you to give us more information than the police of Gotham City (apart from the obvious fact that you can use more than one character messages).

You can choose to contact PDC in several different ways:

- 1. E-mail support@pdc.kth.se.**
- 2. Call PDC Support at +46-8-7907800**
- 3. Visit us at Teknikringen 14, floor (call ahead).**

We do our best to answer your questions on work days during office hours. For most errands e-mail is the quickest way to get help.

In total we are less than a handful people running PDC Support, answering your e-mails, installing software, creating accounts and AFS-volumes plus keeping Lenngren, Lucidor, Key, Ferlin and others up and running.

When PDC Support has answered an e-mail from you, the subject of that e-mail will contain a five-digit number in brackets: [RT# VWXYZ]. This is your request ticket (RT) number. When you send a new e-mail to us and the subject of that e-mail contains the same RT number, PDC Support will automatically see the history of that communication!

Please note that PDC Support does not send you any automatic e-mail replies. Behind every answer is a person. This is why response times vary and why it might take some time before you receive your RT number when you have sent us a new support question.

Batman has his personal man servant, Alfred, who handles incoming calls and seems to run all the gadgets and contraptions found in the Bat cave.

At PDC we take turns answering incoming support e-mails and together take care of every day operational tasks such as replacing hardware, updating software and preparing for future systems.

An advantage of this for you, the user, is that if you send e-mails to support@pdc.kth.se, all of us will be able to read your question--and answer it. If you had contacted the single person Alfred and he was unavailable, your message would have had no chance of being answered.

At PDC we are not superheroes like Batman, but hopefully not villains either! We have a common problem: PDC users don't take vacations, holi-

cern.ch/castor/] and LHC Era Monitoring (Lemon) [<http://lemon.web.cern.ch/lemon/index.shtml>] as well as giving a storage administrative support to LHC [<http://lh.web.cern.ch/lhc/>] experiments inside the Worldwide LHC Computing Grid (WLCG) [<http://lwg.web.cern.ch/LCG/>]. He started his working experience at PDC in March. "Vast" describes the TIER1 storage he has managed", says Director Erwin Laure, "and also the experience he brings to help further develop storage technologies at PDC."

Staff Focus



THOMAS ZANGERL

Thomas Zangerl started working at PDC on March 9th as an NDGF-funded ARC developer. That means that his work will mainly focus on improving the ARC Grid middleware [<http://www.nordugrid.org/middleware/>] within the context of the NDGF project [<http://www.ndgf.org/>]. Organizationally, PDC supplies his working place and the contract but, NDGF cover his cost to PDC. On a more personal note, Thomas is from Austria and studied at the University of Innsbruck, where he graduated with an MSc in computer science in February. Thomas says he is looking forward to working at PDC and living in the beautiful city of Stockholm. Director Erwin Laure says "One definition of an arc is a continuous progression or line of development. Thomas will nicely complement our arc in Grid activity here at PDC."

PDC Related Events (Sponsored/Associated)

PDC's Erwin Laure is Demo Chair at HPDC 2009.

HPDC 2009 – INTERNATIONAL SYMPOSIUM ON HIGH PERFORMANCE DISTRIBUTED COMPUTING

June 11-13, 2009, Munich, Germany. Info: <http://www.lrz-muenchen.de/hpdc2009/>

Check out PDC computers on the latest Top500 list at ISC'09.

INTERNATIONAL SUPERCOMPUTING CONFERENCE ISC'09

June 23-26, 2009, Hamburg, Germany. Info: <http://www.supercomp.de/iscog/>

Learn about Grid and hear PDC's Åke Edlund give a lecture on clouds computing at ISSGCo9.

ISSGCO9 – INTERNATIONAL SUMMER SCHOOL ON GRID COMPUTING 2009

July 5-17, 2009, Nice, France. Info: <http://www.iceage-eu.org/issgco9/index.cfm>

Learn about HPC and PDC at the annual summer school. This is a course within the National Graduate School in Scientific Computing (NGSSC) and is also endorsed by the KCSE Graduate School.

PDC SUMMER SCHOOL 2009

INTRODUCTION TO HIGH-PERFORMANCE COMPUTING

August 17-28, 2009, KTH main Campus, Stockholm, Sweden

<http://www.pdc.kth.se/education/summer-school>

...from page 5

days or sleep - PDC Support does! During every weekend, Christmas, Easter and parts of the summer, PDC Support is closed. If something breaks during this time it might stay broken. Therefore it is important that users plan ahead; register to PDC's mailing lists and check our Flash News and web page for the latest information.

On any given day we have approximately 20 open support questions. But, at times (when something is broken or during vacations) e-mails pile up and can be close to a hundred messages!

At times like that, the best way for you as a user to help us, is to try to keep your confidence in PDC Support. We do remember your requests! We do not delete them and we will answer everyone.

PDC Support is currently making an effort to improve our information to, and communication with, our users. The PDC Web page is being updated and a user forum will be started. We are also going to introduce support meetings and other get-togethers. Keep an eye open in coming PDC Newsletters for more information!

* **Sign up on PDC's mailing lists:** www.pdc.kth.se/about/ mailing-lists

* **Check out Flash News:** www.pdc.kth.se/cgi-bin/flash/flash.py

* **Read more about PDC Support at** www.pdc.kth.se

Sources

We can recommend the following sources for other interesting HPC opportunities and events:

PRACE

<http://www.prace-project.eu/prototype-access>

<http://www.prace-project.eu/hpc-training-events>

DEISA

<http://www.deisa.eu/science/decide>

http://www.deisa.eu/news_press/symposium/Amsterdam2009/deisa-symposium-amsterdam-may-11-13-2009

TERAGRID

<http://www.teragrid.org/tgog/>

US DEPARTMENT OF ENERGY

<http://hpc.science.doe.gov/>

HPC UNIVERSITY

<http://www.hpcuniv.org/events/current/>

HPCWIRE

<http://www.hpcwire.com/events/>

NETLIB

<http://www.netlib.org/confdb/>

CERN COURIER

<http://cerncourier.com/cws/events>

EGEE

<http://events.eu-eg ee.org/>

LINUX JOURNAL

<http://www.linuxjournal.com/events>