

PDC Center for High Performance Computing

# PDC Newsletter

**Ninth Annual SeRC Meeting** - Olivia Eriksson, SeRC, page 3

**WASP Extended to AI** - Anders Ynnerman, Linköping University, page 4

**Introduction to PDC Course** - Thor Wikfeldt, PDC, page 5

**The Nordic Research Software Engineer Initiative** - Thor Wikfeldt, PDC, page 6

**Software Carpentry in Stockholm** - Thor Wikfeldt, PDC, page 7

**Overview of the First Year of HPCE3 Project Activities** - Lilit Axner, PDC, page 8

**Escalating Exascale** - Stefano Markidis, CST, page 9

**Pub Revives Retrogames** - Gert Svensson & Henric Zazzi, PDC, page 9

**First PDC Industry Day Great Success** - Lilit Axner, PDC, page 10

**PDC-SeRC Seminar Series** - Thor Wikfeldt, PDC, page 11

**Cross-Infrastructure Services: Towards Seamless Access** - Michaela Barth, PDC, page 12

**Next Phase: PRACE-6IP** - Lilit Axner, PDC, page 15

**BioExcel Update** - Rossen Apostolov, PDC, page 16





**Erwin Laure**  
**Director, PDC**

### **The PDC Newsletter is published by PDC at CSC, KTH.**

PDC operates leading-edge, high performance computers as easily-accessible national resources. These resources are primarily available for Swedish academic research and education. PDC, which is hosted by CSC, KTH, is one of the six centres in the Swedish National Infrastructure for Computing (SNIC).

**Editor:** Erwin Laure

**Editorial Staff:** Genet Edmondson

**Layout:** Genet Edmondson

**Graphic Design:** Maria Malmqvist

**E-mail:** [support@pdc.kth.se](mailto:support@pdc.kth.se)

ISSN 1401-9671

### **Cover**

One aspect of WASP research works on small autonomous drones participating in a collaboration. The drones are equipped with technology for object tracking and can, for example, detect and follow people in need of help from far away. Photo credit: Göran Billeon

### **How to Contact PDC**

**Visiting address:**

Teknikringen 14, "Plan 4",  
KTH, Stockholm

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

**E-mail:** [support@pdc.kth.se](mailto:support@pdc.kth.se)

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

**Phone:** +46 8 790 7800

## **Editorial**

Artificial Intelligence and Machine Learning are the new buzzwords one can hardly escape. But what is behind them and what impact will this technology have on our lives? The Wallenberg Autonomous Systems and Software Program (WASP) has recently been extended towards AI with an additional generous grant from the Knut and Alice Wallenberg Foundation (KAW). WASP will advance AI technologies and their applications, as you can see in our cover article.

This spring PDC has continued its collaboration with SeRC, the Swedish e-Science Research Centre. During the SeRC annual meeting in May this year, plans for additional multidisciplinary collaboration projects were refined further. Meanwhile the joint PDC-SeRC seminar series is showcasing the inspiring ongoing work at SeRC.

Another interdisciplinary collaboration program, the European HPC-Europa3 program (HPCE3), started successfully this spring. HPCE3 is facilitating research visits by computational scientists to KTH and the Stockholm region, and the first 14 visitors have been hosted by PDC. Check out the breadth of their research areas in this newsletter.

Writing good software is a challenge, particularly in academic settings. For some time now, we have been training people in software engineering techniques as part of the work of the Nordic e-Infrastructure Collaboration (NeIC). Enthusiastic members are now trying to form a larger network, The Nordic Research Software Engineer Initiative – check out their activities and consider joining!

A lot is happening on the European e-infrastructure front as well. PRACE is planning for the 6th phase of its implementation projects and the European Open Science Cloud (EOSC) is shaping up through projects like the EOSC-hub.

Finally, together with researchers from the CST department at KTH, we are continuing our efforts towards Exascale computing and a couple of new projects will start in the autumn. These will contribute to the bold European vision of building European Exascale systems, an endeavour that is currently being formalized through the EuroHPC Joint Undertaking, which is expected to be approved by the European Council this autumn. You will hear more about this in future editions of this newsletter.

With this, I wish you a relaxing summer - but don't forget: PDC's systems do not have a summer break. Keep them busy computing while you enjoy the beach!

*Erwin Laure, Director PDC*

# Ninth Annual SeRC Meeting

Olivia Eriksson, SeRC

The ninth annual meeting of the Swedish e-Science Research Centre (SeRC) was held on the 14th and 15th of May 2018. The meeting took place at the beautiful [Villa Aske](#). As the villa is in the countryside north-west of Stockholm, buses were chartered to transport participants between KTH and the venue.

Along with presentations about seven proposals for Multidisciplinary Collaboration Programmes that SeRC may support, the program for the meeting featured presentations by SeRC researchers, as well as external speakers:

- *Ericsson and e-Science* by Kristina Gold, Ericsson,
- *OpenSpace* by Anders Ynnerman, SeRC,
- *Computing in Science Education* by Anders Malthe-Sorensen, Center for Computing in Science Education, University of Oslo,
- *ELIXIR and e-Science* by Niklas Blomberg, ELIXIR,
- *The search for cell assemblies in the brain* by Arvind Kumar, Brain-IT, SeRC,
- *Data driven Material Science* by Matthias Scheffler, Fritz Haber Institute,
- *Deep learning for prostate cancer histopathology classification* by Mattias Rantalainen, eCPC, SeRC,
- *Unsteady aerodynamics simulations* by Prabal S Negi, SeRC,
- *EuroHPC* by Erik Lindahl, Co-Director, SeRC, and
- *The WASP AI initiative* by Anders Ynnerman, SeRC.

A SeRC board meeting and an Advisory board meeting were held in conjunction with the annual meeting, which was well attended with about a hundred people present, most of whom were SeRC researchers from KTH, Stockholm University, Linköping University and the Karolinska Institute. The plan is that next year's annual SeRC meeting will be held in the spring of 2019 at another interesting location, so keep an eye on <http://e-science.se/news-and-events/events> for details about attending the 10th Annual SeRC Meeting!



Above: SeRC Annual Meeting, 14-15 May 2019, Villa Aske, Bro, Sweden

## In This Issue

### Editorial

Erwin Laure ..... 2

### Staff Focus

Raghunath Vairamuthu ..... 7

Tony Lundgren ..... 10

Xin Li ..... 11

Adam Peplinski ..... 12

Muhammed Ahad ..... 15

### Ninth Annual SeRC Meeting

Olivia Eriksson ..... 3

### WASP Extended to AI

Anders Ynnerman ..... 4

### Introduction to PDC Course

Thor Wikfeldt ..... 5

### The Nordic RSE Initiative

Thor Wikfeldt ..... 6

### Software Carpentry in Stockholm

Thor Wikfeldt ..... 7

### Overview of the First Year of HPCE3 Project Activities

Lilith Axner ..... 8

### Escalating Exascale

Stefano Markidis ..... 9

### PUB Revives Retrogames

Gert Svensson & Henric Zazzi .... 9

### First PDC Industry Day

#### Great Success

Lilith Axner ..... 10

### PDC-SeRC Seminar Series

Thor Wikfeldt ..... 11

### Cross-Infrastructure Services

Michaela Barth ..... 12

### Next Phase: PRACE-6IP

Lilith Axner ..... 15

### BioExcel Update

Rossen Apostolov ..... 16

PDC-Related Events ..... 16

HPC Sources ..... 16

# WASP Extended to AI

Anders Ynnerman, Linköping University

WASP, the Wallenberg Autonomous Systems and Software Program, was launched in 2015 to advance Sweden into a leading position internationally in the area of autonomous systems and their associated software. Autonomous systems are systems that perform advanced tasks and, in many cases, interact with humans, adapting to their environment through information derived from sensors or stored knowledge. For example, robots can be autonomous systems, as can various software applications. Many autonomous systems can also form intelligent systems-of-systems. An example in WASP is the research on collaborative unmanned aerial vehicles (UAVs) forming a fleet of drones with a common mission. Since software is what drives autonomous systems and determines their capabilities, the study of software is crucial in the research and development processes for autonomous systems.

The WASP program was designed to facilitate world-leading academic research on autonomous systems, as well as providing state-of-the-art graduate education in the area, and fostering fruitful interactions between academic

researchers in Sweden and Sweden's leading technology companies. With initial funding of 2 billion SEK, the program has been the largest individual research program in Swedish history right from the beginning! It involves researchers in Electrical Engineering, Computer Engineering, and Computer Science at Sweden's five major Information and Communication Technology (ICT) universities: Chalmers University of Technology, the KTH Royal Institute of Technology, Linköping University, Lund University, and Umeå University.

During the first couple of years of the project, WASP focused on strengthening, expanding, and renewing Swedish competence in autonomous systems and software through funding the strategic recruitment of outstanding leading and young researchers in the field, and by establishing a challenging research program and a national graduate school, and initiating research and development collaborations with Swedish companies.

2017 was the centenary year for the Knut and Alice Wallenberg Foundation (KAW), which provided the funding to establish WASP in 2015. The foundation has a long-standing tradition of contributing towards fundamental research to



Above: One example of AI research in WASP is the research on collaborative unmanned aerial vehicles (UAVs) forming a fleet of drones with a common mission. Here the next test is being prepared in the AI lab at Linköping University. RMAX helicopters will map the terrain while collaborating with each other and with humans.

support the development of new technologies that will assist the long-term development of Sweden's national research and industrial concerns. In line with that tradition, and honouring the centenary of KAW, the foundation made a jubilee donation committing to a total of 1.6 billion SEK over ten years to support research in two of the most important technologies for the future: quantum technology and artificial intelligence (AI).

As a result of this generous donation, the WASP program was extended (in late 2017) to being the Wallenberg AI, Autonomous Systems and Software Program (still known as WASP) and granted 1 billion SEK to build up competence in various aspects of AI – such as machine learning, the underlying mathematics needed for AI software, and deep learning (which is a technique for implementing machine learning that uses multi-layered artificial neural networks to deliver state-of-the-art accuracy in a large variety of tasks like object detection, speech recognition, and language translation, with more applications appearing continuously). The universities in the program (Chalmers, KTH, Linköping, Lund and Umeå), along with Swedish companies, will also contribute towards the funding for the WASP research.

With this additional funding, WASP will man the frontier of AI research and development in Sweden. The grant will also make it possible to develop new graduate schools in AI and recruit senior AI researchers to Sweden. In addition, WASP has reserved 70 million SEK to build a computational infrastructure that will benefit the whole field of AI in Sweden.

The AI research within WASP will be led by KTH researchers, with Professor Danica Kragic, from the School of Electrical Engineering and Computer Science (EECS) at KTH, steering the research on machine learning and next generation AI, while Professor Johan Håstad, also from EECS, will be in charge of the research on mathematical principles underlying AI. During the spring, one of the challenges has been to formulate how the AI expansion within the program will influence the overall research areas of WASP.

“During recent years, the developments in AI based on self-learning systems (such as the advent of autonomous cars, Google translate and medical expert systems) have made the world sit up and take notice. The grant from the Knut and Alice Wallenberg Foundation gives Sweden a fantastic opportunity to develop our fundamental understanding of such new systems, which is vital in order to further the development and uptake of these new technologies by Swedish society,” says Lars Nielsen, the Program Director for the Wallenberg AI, Autonomous Systems and Software Program.

---

## Introduction to PDC Course

Thor Wikfeldt, PDC

The course “Introduction to PDC Systems” is held biannually and serves to introduce researchers to the fundamentals of high performance computing (HPC), while focusing particularly on the HPC infrastructure at PDC. The topics that are covered in the course include: background information about HPC and PDC, information about the computer clusters available at PDC, how to apply to use PDC resources, how to get a PDC account and log in to PDC clusters, how to run software on different nodes at PDC, how to efficiently store your research data, and how to compile code at PDC.

The first introductory course for 2018 was held at PDC on the 14th of February and was attended by 15 researchers. The slides from the course are available at <https://github.com/PDC-support/introduction-to-pdc/tree/intro-course>. As well as going through the prepared material, course participants were able to get assistance to set up their laptops for logging in to PDC, and there was a live demonstration of common procedures and commands used when working on the PDC systems, along with several exercises that the participants could solve.

The next introductory course will be held in September or October 2018. So, if using the PDC systems is new to you, it is highly recommended that you attend the course in the autumn since it is likely to significantly speed up the process of becoming comfortable with working in the HPC environment at PDC!

# The Nordic Research Software Engineer Initiative

Thor Wikfeldt, PDC



Scientific software is used by research groups across a wide range of disciplines, and much of that software is developed within the research groups themselves. In most cases Ph.D. students and postdoctoral researchers are tasked with writing the software, although sometimes dedicated scientific programmers are recruited to larger research groups. When it comes to high performance computing (HPC) projects in Sweden, the Swedish National Infrastructure for Computing (SNIC) centres employ application experts who can provide long-term support to the researchers to assist them with developing software and deploying it on HPC resources. However, while the expertise and the demand for advanced application and software support exists, the relevant experts and researchers are often disconnected and unaware of each other. As a result, there are many disparate specialised clusters, with little exchange of technical expertise, and researchers outside the HPC community lack a common entry point to seek professional advice and assistance on writing scientific software.

The environment for research software development around the Nordic countries was discussed in a speed-blogging session at the all-hands meeting of the Nordic e-Infrastructure Collaboration (NeIC) held in early 2018 in Norway. This resulted in a [published blog post](#) where various challenges faced by research software developers were identified, and possible ways to improve them were outlined. Following up on these discussions, a group of enthusiasts gathered in a conference centre outside Stockholm and decided to launch a Nordic Research Software Engineer (RSE) network, which was inspired by similar networks that have been established and are presently thriving in the UK, the Netherlands, Canada and Germany.

But what is an RSE, and what is the point of creating a network to connect people doing RSE work across the Nordic countries? RSE stands for Research Software Engineer, a term chosen after careful deliberation in the UK (where the first RSE network was started) to marry aspects of software engineering with academic research. Some RSEs start off as researchers who spend time developing software to make progress with their research. Because they enjoy this work and have invested time and effort in developing specialist skills, these researchers continue to focus on software and its use in research. Others start off from a more conventional software development background and are drawn to research by the challenge of using software to contribute to research. This combination of skills is extremely valuable for science, but RSEs often lack a formal place in the academic system. This means there is no easy way to recognize their contribution, to reward them, or to represent their views. Without a name, it is difficult for people to rally around a cause, so the term Research Software Engineer was created. The Nordic-RSE network will now start working to raise awareness of the importance of this role for research and to bring the RSE community together across the Nordic countries by organizing conferences, workshops and online forums.

As a first step towards building this community, the Nordic RSE initiative is conducting a short five minute survey on the situation regarding research software development in the Nordic countries. If you write research software, then this survey is for you! To participate in the survey, please follow the link: <http://nordic-rse.org/survey>. And if you have colleagues or friends who write research software, please forward this information to them! To be informed about the results of this survey and other items of interest to research software developers, visit <http://nordic-rse.org>, follow the initiative on [Twitter](#) and/or sign up for the [Nordic RSE mailing list](#). Finally, if you would like to join us in building the community, please read <http://nordic-rse.org/join>.

Note: This article is based on text from <http://nordic-rse.org> and <https://neic.no/news/2018/05/04/building-a-community>.

# Software Carpentry in Stockholm

Thor Wikfeldt, PDC

The [Software Carpentry Foundation](#) is a non-profit organization devoted to teaching researchers the computing skills they need to get more done in less time and with less pain. After its start in 1998 as a week-long training course at US national laboratories, Software Carpentry has evolved into a worldwide volunteer movement - hundreds of workshops have been held around the world and the number of attendants has surpassed 34,000. After the success of Software Carpentry, similar initiatives were launched to teach other more specific computational skills, including [Data Carpentry](#) for domain-specific data management, [Library Carpentry](#) for computational skills relevant to librarians, and [HPC Carpentry](#) for users of HPC resources. Recently, these initiatives were joined under a common organisational structure called [The Carpentries](#).

With support from the Swedish National Infrastructure for Computing (SNIC), a Software Carpentry workshop was held at the KTH campus on 12-13 March 2018. The event was advertised through the SNIC training newsletter and, judging by the rapid rate at which registrations came in for the workshop, there is clearly a large demand for these types of courses in Stockholm. All Software Carpentry course material is directed towards students and researchers who want to learn the basics of fundamental tools for computational work. The topics chosen for the Stockholm workshop were (1) the Unix shell for automating and speeding up tasks, (2) Git for version control of scripts, programs and plain text files, and (3) the Python programming language.

Software Carpentry workshops are more foundational in nature compared to workshops taught in the related [CodeRefinery project](#) (run by the Nordic e-Infrastructure Collaboration). The latter focus on tools and best practices for scientific software development for researchers who already write code as part of their work and, as such, provide a natural second step after attending a Software Carpentry workshop. A closer collaboration between CodeRefinery and The Carpentries is envisioned in the future, which will hopefully mean that more Carpentry workshops can be held in the Nordic countries.

## Staff Focus



Raghunath Vairamuthu

Raghunath Vairamuthu is a first year Masters student specialising in Machine Learning at KTH. He is taking courses such as Artificial Intelligence, Neuroscience and Deep Learning. His background is Electronics and Communication Engineering in his Bachelor's degree and he worked on designing NVIDIA GPUs for a few years after graduating. At present, Raghunath is working part-time in first-line support at PDC, taking care of account management, addressing support requests, and documenting helpful resources.

Raghunath works on hobby robotics projects during his free time and likes playing and creating video games. He also enjoys playing football and is a member of the IM association which helps the young student community and assists with the welfare of people who have newly arrived in Sweden.



Above: A few of the happy participants at the Stockholm Software Carpentry workshop, 12-13 March 2018

# Overview of the First Year of HPCE3 Project Activities

Lilit Axner, PDC

As you already know from previous editions of our newsletter, the HPC-Europa3 (or HPCE3) project is based on a programme of visits, in the form of traditional transnational access, with researchers visiting high performance computing (HPC) centres and/or scientific hosts. The visitors are funded for travel, accommodation and subsistence, and provided with an amount of computing time suitable for their approved project.

The 1st of May this year marked the first anniversary of the HPCE3 project, and in the year since the project started in May 2017, HPCE3 has already launched four calls for applications.

During the first three of those four calls (The evaluation of the fourth call is on-going at the time of writing.) HPCE3 received  $70+23+70 = 163$  applications, of which  $56+20+47 = 123$  have been accepted for visits to one of the nine partner countries. Of the 123 applications that have been accepted, 14 have been applications for researchers to visit PDC and use computing time on Beskow and Tegner, while having a collaborative host researcher at KTH, Stockholm University or Uppsala University. In this first year alone, PDC has already hosted 35% of the visitors that it committed to hosting during the four years of the HPCE3 project.


The following table shows details about the researchers who have visited PDC so far – including the visitor's home country and university, their research discipline, the software they have used, and the university or company in Sweden with which they were collaborating.

Country	Home University	Scientific Discipline	University or Company in Sweden	Software
Lithuania	Center for Physical Sciences and Technology	Material Sciences	KTH	VASP
Latvia	Riga Technical University	Structural Mechanics	KTH	own software
Slovenia	University of Ljubljana	Fusion Plasma Physics	KTH	SMARDDA
Greece	Technological Educational Institution TEI of West Macedonia	Environmental and Pollution Control Engineering	Stockholm University	Gaussian
Denmark	Niels Bohr Institute, University of Copenhagen	Theoretical Astrophysics	Stockholm University	PENCIL
Serbia	Institute for Multidisciplinary Research, University of Belgrade	Material Sciences	KTH	Quantum ESPRESSO, Crystal
Switzerland	EPFL	Computer Science	Airinnova AB	SU2
Belarus	Belarusian State University of Informatics and Radioelectronics	Material Sciences	KTH	VASP
Lithuania	Vilnius University	Chemical Physics	Stockholm University	Dalton
Lithuania	Vilnius University	Chemical Physics	Stockholm University	Dalton
Serbia	University of Belgrade	Material Sciences	KTH	Quantum ESPRESSO, PWscf
UK	University of Cambridge	Chemistry	Uppsala University	CP2K
Latvia	University of Latvia	Material Sciences	KTH	CRYSTAL17
Latvia	Latvian Institute of Organic Synthesis	Life Sciences	KTH	GROMACS





Above: Prof. Natalia Skorodumova (KTH) - HPCE3 host  
 Below (left to right): Nikolas Benetis (Greece), Leon Kos (Slovenia) & Janis Sliseris (Latvia), HPCE3 visitors at PDC



## Reminder re applying to HPCE3

**PDC has committed to hosting about 40 visitors from all over the world, with priority being given to visitors from the Baltic countries (to tighten the collaboration between Sweden and the Baltic countries, as well as promoting the use of HPC in the Baltic countries). HPCE3 calls for applications are always open, however there are four cut-offs per year. Currently HPCE3 is accepting applications for the fifth cut-off that will end on the 13th of September 2018.**

## PDC Survey 2018

Thank you to everyone who took the time to participate in the recent survey run by PDC! Your input was much appreciated – it will lead PDC to increase the support we offer as it enables us evaluate what our customers think is lacking. The survey results will also increase our awareness regarding the additional hardware that is likely to be needed in the future to support the research community in Sweden.

The information from the survey is currently being evaluated and the results of the survey will be published in the beginning of the autumn.

## Escalating Exascale

Stephano Markidis, CST

September 2018 is an important crossover point for European exascale projects at PDC. Three existing projects – SAGE, INTERTWinE and AllScale – will hand over their exascale research to three new projects: EPiGRAM-HS, Sage2 and VESTEC. EPiGRAM-HS and Sage2 will build upon the work of two existing projects that PDC has contributed to (EPiGRAM and SAGE). PDC will lead the EPiGRAM-HS (Exascale Programming Model for Heterogeneous systems) project. As in the first EPiGRAM project, the research topic will be the design and development of programming interfaces, in particular MPI and GASPI, towards exascale. However, the new EPiGRAM-HS project specifically targets distributed systems that are heterogeneous (both in terms of their compute and memory capacity). Sage2 will continue the work started in SAGE, still focusing on object storage technologies and data-centric computing. While EPiGRAM-HS and Sage2 build on previous work, VESTEC is a new exascale project for the visualization and analysis of large data volumes on exascale supercomputers. All three projects will start with kick-off meetings in September 2018; stay tuned for further details!

## Pub Revives Retrogames

Gert Svensson & Henric Zazzi, PDC

The PDC Pub provides an informal opportunity for researchers and students at KTH to discuss high performance computing research and related matters with PDC staff while enjoying refreshments. This year the Pub was held on the 9th of May. As usual, Gert Svensson (PDC's Deputy Director) gave guided tours of the PDC computer hall, while a vintage game competition was held in the SeRC room, with a prize for whoever achieved the highest score in Pinball. With a score rumoured to be around 44 million, the lucky winner was awarded a free meal at the Syster O Bror restaurant. People could also try out a selection of different retrogames on the big screen and relive classic titles like Doom and Dig Dug.

## Staff Focus



Tony Lundgren

Tony Lundgren works in first line support at PDC and is just finishing the first year of his Masters degree in embedded software at KTH, where he is studying topics like computer architecture and real-time systems. Tony is a big supporter of free software and the ethics that come with it. Outside of computers, he enjoys movies both about and from earlier decades and lots of different kinds of music.

*Below: PDC Industry Day, Stockholm, 1 March 2018 - discussions (below) participants listening to seminars (bottom left), Torbjörn Larsson from Creo Dynamics (bottom centre) and Mengmeng Zhang from Airinnova (bottom right)*



## First PDC Industry Day Great Success


Lilit Axner, PDC

On the 1st of March 2018 PDC organized an event for small to medium-sized enterprises (SMEs) and large companies that have an interest in high performance computing (HPC). The intention with the event was for the participants to discuss the use of HPC by non-academic researchers, as well as their HPC requirements, and to introduce recent HPC developments and collaborative research possibilities (both on the national and international levels) that enable industrial concerns to work with academic HPC systems, in addition to showcasing successful examples of such research.

Over 30 people attended the event – coming from 13 different enterprises and organizations, such as PDC's strategic partner Scania, as well as Westinghouse Electric Sweden, Sandvik Coromant, the ESI Group and Arm, along with several SMEs including Airinnova, Creo Dynamics, Adaptive Simulations and FS Dynamics.

There were many stimulating talks and plenty of interesting questions were brought to the table. Various companies demonstrated impressive results that were produced thanks to the use of HPC. Airinnova and Creo Dynamics are two of the Swedish SMEs that have been using HPC systems under the auspices of the international projects PRACE and HPCEuropa3. Both of those companies highlighted the importance of having opportunities to access HPC systems - free of charge. This is especially valuable for SMEs as smaller companies rarely have the resources to provide their own HPC systems.

Scania AB, being a strategic partner to PDC, walked people through the whole collaboration process and also displayed some research results that were recently obtained on one of the PDC HPC systems called Beskow. Westinghouse and Sandvik Coromant presented interesting results obtained using their own computational clusters and discussed their future HPC plans, along with discussing their needs to adopt even more extensive use of HPC. Overall the participants evaluated the event as having been quite successful and noted particularly that it had been a useful forum for sharing information between the different enterprises about their HPC requirements and how they use HPC.



## PDC-SeRC Seminar Series

Thor Wikfeldt, PDC

In the autumn last year, PDC and the Swedish e-Science Research Centre (SeRC) started a new series of inspiring seminars about cutting-edge current e-Science research. The seminars are held roughly once per month (in the SeRC Open Space at PDC) at lunchtime. A light lunch with hot drinks is offered to attendees, which creates a convivial atmosphere where people are welcome to ask questions and discuss the research.

The seminar series continued during the first half of 2018 with a further five seminars:

- “Spectroscopy of molecules in their biological environment”, by Prof. Mathieu Linares, associate professor in the Department of Theoretical Chemistry and Biology at KTH (24 January),
- “GUI frameworks for fusion codes SOLPS-ITER and SMITER” by Dr. Leon Kos, researcher at the Faculty of Mechanical Engineering at the University of Ljubljana, Slovenia (14 February),
- “Scalable simulation of cortical associative memory” by Dr. Anders Lansner, professor at the Department of Computational Science and Technology, KTH, and also at Stockholm University (15 March),
- “Parallelization strategies in GROMACS”, Berk Hess, project leader for GROMACS development at KTH (25 April), and
- “Software Development for Quantum Molecular Modelling on PDC Cluster Resources” by Patrick Norman, Head of Department and professor at the Department of Theoretical Chemistry and Biology at KTH (16 May).

The seminar series has been well received by people in the e-Science research communities, and consequently PDC, together with SeRC, plans to continue with these monthly seminars after the summer break: inviting outstanding researchers from various disciplines to present exciting results from the forefront of their research fields. The seminars are growing into an established tradition in Swedish e-Science and contribute to inter-disciplinary discussions and cross-fertilization of ideas.

We invite you to attend the seminars when they start again in September 2018. All the seminars will be advertised on the PDC website and through SeRC mailing lists, but if you would like to receive personal email announcements, please contact Thor Wikfeldt ([kthw@kth.se](mailto:kthw@kth.se)), who organizes the seminars at PDC. Finally, we also welcome volunteer speakers, so please get in touch with Thor if you are interested in giving a PDC-SeRC seminar!

## Staff Focus



Xin Li

Xin Li obtained his Ph.D. in the department of Theoretical Chemistry and Biology, KTH. During his Ph.D., Xin focused on multiscale modelling of optical and magnetic properties of functional molecules that may find applications in biosensing and bioimaging. He also developed tools for the optimization of force fields that are used in molecular dynamics simulations. After his Ph.D., Xin continued his research as a postdoc at KTH, and later at Stanford University, with a focus on composite organic-inorganic systems and multichromophoric systems, for which he developed capacitance-polarizability force field and ab initio exciton models for efficient simulation and large-scale parallelization.

Xin is currently working part-time as an application expert at PDC. He provides support for scientific software users, and works as a developer of quantum chemistry code moving towards exascale computing. He is the main developer of the Exciton code, and co-developer of the Fiesta and VeloxChem codes.

In his free time Xin likes to play badminton, soccer, and floorball. He also enjoys watching movies and reading novels, in particular those based on science fiction.



## Staff Focus



Adam Peplinski

Adam Peplinski's background is in theoretical physics, which he studied at Nicolaus Copernicus University in Poland. Later he became interested in numerical simulations of astrophysical phenomena, so Adam started Ph.D. studies at Stockholm University working on the interaction between forming planets with the surrounding circumstellar disc. During that time, he also gained some experience in system administration, building a small Beowulf cluster to perform his computations. Later Adam joined KTH Mechanics and PDC as an application expert and now works closely with the flow stability/turbulence transition group. His main interests are high order methods for computational fluid dynamics (CFD) and their optimisation using mesh adaptivity; within the European projects CRESTA and ExaFLOW, he implemented h-type Adaptive Mesh Refinement in the spectral element method (SEM) code Nek5000. In his free time, Adam likes to hike and paraglide.

## Cross-Infrastructure Services: Towards Seamless Access

Michaela Barth, PDC

EUDAT is dead, long live EUDAT!

From 2011 till early 2018, the work of EUDAT was supported by two European Commission (EC) grants: the first under the Seventh Framework Programme (FP7), where the project was known as EUDAT, and the second under the Horizon 2020 Programme (H2020), with the project being referred to as EUDAT2020. Both of these EUDAT projects have worked on developing a pan-European Collaborative Data Infrastructure (CDI) for handling European research data.

This spring the EUDAT2020 project ended, but the EUDAT Collaborative Data Infrastructure and the consolidation of its services will continue: the EUDAT CDI is led by a council composed of service provider representatives and formally represented by EUDAT Ltd – a limited liability company established and registered under Finnish law at the end of February 2018. EUDAT Ltd is hosting the administrative functions (management, operations and technical coordination, plus user engagement and outreach) of the CDI agreement and will represent the CDI network at the European level.

This article describes PDC's main contributions to EUDAT2020 and discusses where and how the outcomes of that work will be used in the future.

Within EUDAT2020, the Swedish National Infrastructure for Computing (SNIC) was represented by PDC. As mentioned in earlier newsletters (2014 #2 & 2015 #1), PDC made a significant and major contribution towards the development of the B2SHARE service (which enables researchers to reliably store and share small-scale research data). PDC was also responsible for coordinating the B2HANDLE operations (which provide unique persistent identifiers – often referred to as PIDs or handles – for data items stored via EUDAT) and several major upgrades and configuration changes to the service. These modifications were undertaken in order to make the service more secure, robust and scalable. Both of these services were, of course, also operated and available to numerous national and international research groups. Alongside – and in parallel to – this work, the Swedish National Storage Infrastructure has been federated and now uses the EUDAT B2SAFE service, which provides a bridge between the Swedish national and European e-infrastructures.

EUDAT was also aiming to harmonize access to both data and compute services. Joint pilot projects combining HPC and data

services were conducted with [PRACE](#), and PDC supported some of those projects, such as the CHARTERED data pilot within material sciences, and its follow-up CHARTERED2, that modelled CHARGE Transfer dynamics by applying time-dependent Density functional theory. Similarly, EUDAT collaborated with [EGI](#) to enable joint access to Data, HTC and Cloud Computing Resources.

Initially, candidate research communities were identified. Typically these were relevant European research infrastructures that were already collaborating with one, or preferably both, of the EUDAT and EGI infrastructures and coming from the fields of earth sciences (such as EPOS and ICOS), bioinformatics (for example, BBMRI and ELIXIR) and space physics (like EISCAT-3D.) Then the primary requirements of those candidate research infrastructures for connecting EGI and EUDAT services into a joint cross-infrastructure offering (with perceived seamless access) were collected.

This information was used to define a generic use case, which was a major milestone for EGI and EUDAT in terms of understanding each other's interfaces, as well as the desired and necessary base functionality. It was not long after this that the first working demonstration accessing both infrastructures with a single user identity to stage data between EGI Federated Cloud services and EUDAT services could be presented, which successfully demonstrated the principle concepts of the intended interoperability. In the second year of the EUDAT EGI pilot activity, three early adopter use cases (from ICOS, EPOS and ENES) were selected from the most mature candidates identified, and they were then supported in developing their workflows and the data streams within those workflows. At the same time, these early adopters provided very useful feedback to EGI and EUDAT in regard to building infrastructure services for the future. This activity was deemed to be very useful by all parties and was thus extended for a third year.

In addition, PDC made contributions to work on harmonizing access policies which again provided the basis for design choices – this time



within the [AARC](#) blueprint architecture and also when defining the pilot studies within the AARC-2 project that deal with EGI-EUDAT AAI interoperability – these

pilots will continue to focus on aspects of the AAI interoperability.

The work done in this task also finds a logical continuation in the Competence Centres within the [EOSC-hub project](#), which will also guarantee future interoperability between EUDAT's B2ACCESS and EGI's Check-in (which is a proxy-service that enables access to EGI services and resources via federated authentication). [ICOS](#), one of the research communities that PDC worked very closely with, has set up an ICOS Competence Centre under the EOSC-hub project.



The EOSC-hub project started in January 2018 (with an initial duration of three years) as one of the first steps towards the realisation of the vision of a European Open Science Cloud (EOSC) by the European Commission. EOSC-hub will provide a horizontal pan-European consolidated technical infrastructure ecosystem for such an EOSC, based on the building blocks provided by EGI, EUDAT, PRACE and the [INDIGO-DataCloud](#). Here a horizontal ecosystem means one that the users perceive as being a single infrastructure that spans everything that is there – in other words, with just one login, each user will be able to access all the different infrastructure services that he or she needs to do his/her research. The goal is to produce a federated integration and management system for the future EOSC, with some centralized core services (like the joint catalogue of resources and services, for further details see the [beta version](#)), and to emphasise the federated operation of services under an open community-led framework.



In parallel, the **EOSCpilot project**, which started in January 2017 and is supposed to run for two years, supports the development of the first phase of the EOSC through an increasing number (currently 15) of science demonstrators. The whole process of defining the exact EOSC governance framework and the EOSC rules of engagement is intended to be stakeholder-driven, evidence-based and built on community consent.

Many people are not yet completely clear what the EOSC is really about: is it just a new funding model, or is it an easy access channel to a catalogue of services? Even the name might be confusing to some, since cloud computing per se is not a central aspect of the EOSC.

While cloud computing may be one of the spectrum of services provided by the EOSC, the word “cloud” should be interpreted in a broader sense. The emphasis is really on Open Science and Open Data – where science and research is put into the driving seat of e-infrastructure provisioning for advanced data-driven research and where researchers are given the right means to handle research data in a way that supports the FAIR (findable, accessible, interoperable and reusable) principles.

The EOSC vision is best explained by taking the perspective of the modern researcher: researchers want data handling services that are 100% reliable and that answer the researchers’ needs without the researchers even having to think about the complexities of the underlying system (while still making automation and complete control possible – through clearly defined APIs – for the researchers who want to be able to automatize processes to suit their needs). The services should be agile, reacting quickly to requests, and (like the research itself) the services should be continuously adapting. It should also be very easy for people to start using the EOSC services for their research. However, researchers will need to do more than just play the role of users of the EOSC; since many research infrastructures are

already providing specialised data handling services themselves, at least some research communities are going to have to take on part of the responsibility for the provision and co-development of suitable research e-infrastructure services.

Most researchers probably do not see themselves as data-producing factories, yet they really are the driving force in the data-driven research process, and hence they are also going to have to play a role in the quality assurance for data handling services, for example, by reviewing services, participating in the co-development of services and channelling some of their research funds into developing and supporting such services. Yes, money needs to be spent not just on doing research but also on services for handling the resulting research data! Since the EOSC will not only be a trusted virtual environment with free open and seamless services for data storage, management, analysis, sharing and re-use of data across research disciplines with well-established and defined access channels, it will really be a new type of funding model, too. All the power will be in the hand of the researchers; they will be the ones regulating the funding stream coming from the European Commission to the service providers.

The EOSC will also include a concept of virtual research environments to help research communities to find solutions themselves and generally foster interaction between thematic facilities like data infrastructures and scientific clouds, thereby



*Above: Dr. Michaela Barth being interviewed on collaboration enabling, as well as connecting, services and user communities at the EUDAT 2018 conference: “Putting the EOSC vision into practice”, Porto, Portugal, 22-25 January 2018.*

facilitating interdisciplinary research and effectively bridging today's fragmentation of research into different disciplines to leverage research investments. The EC has just endorsed a new EOSC roadmap: <https://www.eosc-hub.eu/news/eu-competitiveness-council-endorsed-implementation-roadmap-european-open-science-cloud>.



## Next Phase: PRACE-6IP

Lilit Axner, PDC

We are glad to announce that PRACE is moving on and has just submitted the project application for the PRACE 6th Implementation Phase (which would be known as PRACE-6IP) to the European Commission under the H2020 calls. The Swedish National Infrastructure for Computing (SNIC) with three Swedish high performance computing (HPC) centres – namely PDC, the National Supercomputing Centre (NSC) at Linköping University, and the High Performance Computing Center North (HPC2N) at Umeå University – will continue to be one of the 25 PRACE partners.

In general, the activities within PRACE-6IP will not differ much from the activities that are being carried out now within 5IP. However, there two major changes.

1. The DECI optional program will be reorganized and merged even more closely within PRACE. This change will not affect PRACE users/researchers as such but is more of an internal reorganization process.
2. A new work package within PRACE-6IP will be dedicated to so-called “Forward-looking Software Solutions Projects”, where ten software projects will be selected to be worked on. The aim of this is to deliver forward-looking software solutions (to PRACE Tier-0 users and research communities) for the two main challenges induced by the rapidly changing HPC pre-exascale landscape, namely the diversity of hardware (requiring performance portability in codes), and software complexity (requiring separation of concerns in codes).

We expect to get the decision from the European Commission about the PRACE-6IP application during the autumn of 2018. The PRACE research infrastructure is also working intensively to adapt to the new General Data Protection Regulation (GDPR) rules in relation to users' data. In our next newsletter, we hope to inform you about the changes that will have been adopted and put into place for PRACE within the GDPR rules.

## Staff Focus



Muhammed Ahad

Muhammed Ahad is a Masters student in Computer Science at KTH. This is his second Masters; his first was in Scientific Computing. The courses he has studied include Computational Fluid Dynamics, High Performance Computing, and Numerical Analysis. During his first Masters, Muhammed was a summer worker at the company ABB and developed an application for fluid flow simulation based on Smoothed Particle Hydrodynamics (SPH) for the cooling process on a Run Out Table (ROT).

Currently Muhammed is working at PDC in first line support, where his tasks include answering queries from users and writing documentation for PDC resources. He is also working with the INTERTWinE project on the development of GASPI (Global Address Space Programming Interface), a highly scalable, flexible and failure tolerant PGAS (Partitioned Global Address Space) API for asynchronous one-sided communication in high performance computing.

During his free time, Muhammed likes to watch movies, play cricket and travel with his family.

# BioExcel Update

Rossen Apostolov, PDC

BioExcel has been busy during the spring! The BioExcel software engineering teams are passionate about creating effective, usable, high-performance software for life sciences research. They have distilled that knowledge into [guidelines](#) useful for all teams developing software for science.

As many of our applications and tools are being used by industry (in particular in the pharmaceutical branch), BioExcel has begun a series of site visits to companies to share expertise and understand the needs of researchers using our offerings. In March, we visited the [Janssen research and development site](#) in Belgium.

Over the last couple of months, BioExcel has produced a great line up of webinars! We presented the [latest developments of the MARTINI Force Field](#) and a sleek [new web server for structural/dynamical studies of DNA](#). Our industrial (and also academic) users might find the [methods and tools developed by BiKi Technology](#) very interesting and useful, as well as [successful early hit optimization by Novartis](#) using HADDOCK, amongst other tools. And, if you have not tried GROMACS 2018 yet, you really should check out Mark Abraham's overview of its [new features and capabilities](#), which is sure to change your mind! Coming up next is a presentation by BioExcel's collaborators from the [Pistoia Alliance's AbVance project](#).

The 2018 Platform for Advanced Scientific Computing (PASC) conference is being held in Basel in July. It is a great venue for computational scientists in all domains, including the life sciences. BioExcel is running a dedicated [mini-symposium on “Advances and Automation for the Exascale Era”](#) together with our [partners from MolSSI](#). If you are coming to the conference, make sure to join us or drop by our exhibition booth. Another important event is the [BioExcel 2nd SIG Meeting: “Advanced Simulations for Biomolecular Research”](#) during the 17th European Conference on Computational Biology (ECCB 2018). This year's conference is being co-organized with our [partners from Vi-SEEM VRE](#). We have a great line-up of speakers, including a [keynote by Jean-Pierre Changeux](#). So hurry up and send off your abstract for a poster! We are offering a [number of travel bursaries \(up to 500 euro\)](#) to successful candidates!

## PDC-Related Events

### PDC Summer School 2018

13-24 August 2018, KTH, Stockholm

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

## HPC Sources

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

### BioExcel

<http://bioexcel.eu/events>

### CERN

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

<https://home.cern/scientists/events/computing>

### EGI

<https://www.egi.eu/category/events>

### HPC University

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

### HPCwire

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

### NeIC

<http://neic.nordforsk.org>

### PRACE

<http://www.prace-ri.eu/HPC-access>

<http://www.training.prace-ri.eu>

<http://www.prace-ri.eu/events>

<http://www.prace-ri.eu/news>

### SeSE

<http://sese.nu>


### SNIC

<http://www.snic.se/news-events>

<http://docs.snic.se/wiki/Training>

### XSEDE

<https://www.xsede.org>



Above: Pharmaceutical companies use BioExcel software for virtual screening and discovery of novel drug molecules.