



**Concerted action for the European HPC CoEs**

**Project Number: 823964**

**D4.3 – Report on the outcomes of the comprehensive  
European stakeholder workshop on HPC training and  
skills development needs**

**WP4 – Training**



Focus CoE has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement N° 823964

<b>Workpackage:</b>	4	Report on the outcomes of the comprehensive European stakeholder workshop on HPC training and skills development needs
<b>Due Date</b>	30/11/2019	
<b>Author(s):</b>	Simon Wong	ICHEC (NUIG)
	Alan O'Cais	FZJ
<b>Authorized by</b>	Simon Wong	ICHEC (NUIG)
<b>Reviewer</b>	Edouard Audit	CEA
<b>Reviewer</b>	Lilit Axner	KTH
<b>Reviewer</b>	Andrea Townsend-Nicholson	UCL
<b>Final Version Date</b>	26/11/2019	
<b>Dissemination Level</b>	Public	

## List of abbreviations

BDV	Big Data Value Association
CEA	Commissariat à l'énergie atomique et aux énergies alternatives (French Alternative Energy and Atomic Energy Commission)
CERN	Organisation européenne pour la recherche nucléaire (European Organisation for Nuclear Research)
CoE	Centre of Excellence
E-CAM	An e-Infrastructure for Software, Training and Consultancy in Simulation and Modelling
EC	European Commission
ETP4HPC	European Technology Platform for High-Performance Computing
EU	European Union
EU13	[Refer to new EU member states since 2004]
EuroHPC JU	European High Performance Computing Joint Undertaking
EXDCI	European Extreme Data & Computing Initiative
FZJ	Forschungszentrum Jülich (Juelich Research Centre)
HLRS	Höchstleistungsrechenzentrum Stuttgart (High Performance Computing Centre Stuttgart)
HPC	High Performance Computing
HPDA	High Performance Data Analysis
IS-ENES	Infrastructure for the European Network for Earth System modelling
ISV	Independent Software Vendor
JSC	Juelich Supercomputing Centre
KTH	Kungliga Tekniska högskolan (Royal Institute of Technology, Stockholm)
LRZ	Leibniz-Rechenzentrum (Leibniz Supercomputing Center)
LinkSCEEM	Linking Scientific Computing in Europe and the Eastern Mediterranean
MOOC	Massive Open Online Course
MPI	Message Passing Interface
MSc	Master of Science
OpenMP	Open Multi-Processing
POP	Performance Optimisation and Productivity Centre of Excellence in Computing Applications
PRACE	Partnership for Advanced Computing in Europe
SME	Small and Medium-Sized Enterprise
SPOC	Small Private Online Course
SciDAC	Scientific Discovery Through Advanced Computing
UCL	University College London
UK	United Kingdom
US	United States
VI-HPS	Virtual Institute – High Productivity Supercomputing

## Executive Summary

With the vision of the EuroHPC Joint Undertaking (JU) to develop a world-class HPC ecosystem in Europe, training and education has a vital role to achieve its goals by supporting and sustaining a skilled and knowledgeable talent pool. This report describes a workshop that was organised by the FocusCoE project and the European Commission (EC) to bring together experts and stakeholders to define what the training and education requirements are from different domains/target audiences, what the state of the art is as regards HPC training and education, and to examine how the variety of requirements may be met by existing or new European and national programmes on HPC training.

A total of 32 participants attended the workshop, representing the interests of a range of stakeholders and communities such as HPC users, developers, as well as training and education providers. They also include participants with expertise of both academic and industry training needs, as well as representatives from EU13 countries.

For the HPC application user community, there is recognition that building local HPC ecosystems is an essential foundation for training initiatives, especially in European countries where HPC is not well established. Overall the community requires a holistic approach to training and education including where HPC should be introduced at an earlier stage at the undergraduate level, and relevant skills are taught as an intuitive experimental/analytical tool across the entire academic spectrum; some success stories have already been achieved to date but a framework for this to scale will be required. While many CoEs already deliver training that offers a rich mix of face-to-face and online training offerings for HPC application users, the workshop identified future actions that include expansion of existing initiatives and activities, e.g. to target new domains, additional focus on HPDA and mobility programmes that promotes mentoring between application users and developers or experts.

For the HPC application developer community and support staff, there is the ongoing need for developers with in-depth knowledge of both hardware and software, and refresher courses that support lifelong learning. Major topics of discussion include the need to sustain a pool of qualified trainers, who along with application developers require better support for career progression and opportunities. The workshop acknowledged some major training programmes (e.g. at PRACE, national centres) that cater for HPC application developers, but future approaches will have to address the issue of scalability and creating the environment for better appreciation of the role in academia along with prospects for career development opportunities. Discussions also highlighted the need to establish education and traineeship programmes specifically targeted at HPC system administrators.

Across the HPC user and developer communities, the training needs of industry do overlap with those of academia but important differences have been highlighted, e.g. SMEs often require more bespoke training under tighter time scales. The awareness of HPC among industry remains an ongoing issue, and the Fortissimo approach was highlighted as a good example of teaching SMEs to adopt HPC solutions. Future directions for industry-oriented activities include internship programmes that promote cross interaction between academia and industry, and certification for on-the-job training. HPC technology providers from industry also has an important role to play in contributing to HPC training and education.

The discussions surrounding support for career development saw the scope for developing new MSc programmes, or even shorter diploma/certificate programmes, that may incorporate traineeship opportunities with different tracks (e.g. application development, infrastructure). This may be facilitated by interaction and dialogue between universities and HPC training providers. In order to provide more clarity to support HPC related career development, additional effort will

be needed to elucidate the different career paths and professional profiles related to HPC, along with their skills requirements.

In conclusion, the workshop proved to be extremely useful not only to reflect on the current landscape and directions for HPC training, but for the participants themselves who had gained more awareness of national and European initiatives. There is a consensus that the workshop could be the first iteration of future meetings to facilitate communication between the many projects and institutions that conducts training and pave the way to maintain a level of pan-European coordination.

# Table of Contents

1	Introduction	7
2	The objective of the workshop	7
3	Identification of key stakeholders	7
4	Summary of results of the workshop	8
4.1	Identification of needs and problems	8
4.1.1	HPC users	8
4.1.2	HPC application developers and support staff	9
4.1.3	HPC education	9
4.1.4	Supporting HPC-related career paths	9
4.2	Lessons learned	10
4.2.1	HPC Users	10
4.2.2	HPC application developers and support staff	11
4.2.3	HPC education	11
4.2.4	Supporting HPC-related career paths	12
5	Objectives of future actions	13
6	Outlook on future actions	14
7	Conclusion	15
8	Annex	16
8.1	Participants	16
8.2	Workshop agenda	17

# 1 Introduction

HPC training and education plays an important role in sustaining a steady stream of researchers and developers who can exploit the full potential of High Performance Computing (HPC) and modern computational laboratories. Its role is set to grow in the coming years as Europe, via the EuroHPC Joint Undertaking (JU), looks to develop the next generation of technologies, applications and systems towards exascale, as well as enhancing uptake of HPC from industry and achieving excellence in HPC application delivery and use.

As part of the FocusCoE project – which aims to provide support to the European CoEs and assist in the harmonisation of their activities – a workshop was organised in coordination with the European Commission (EC) aimed at gathering HPC training experts and stakeholders both to define what the training requirements are from different domains/target audiences, and to examine how these requirements may be met by existing or new European and national programmes on HPC training.

It is important to highlight that we use the term “training” throughout this document to refer loosely to both training and education. While training is focused on developing skills to accomplish specific tasks, education covers broader aspects of a subject including the theory and knowledge that may not be necessary to carry out the same tasks but there are overlaps between the two in many cases.

This document describes the organisation and the outcome of the “European HPC Training Stakeholders Workshop” that was organised in Brussels on 8<sup>th</sup> October 2019 with 32 participants. It covers the different stages and aspects of the workshop including the objective of the workshop (Section 2), the identification of key stakeholders (Section 3), summary of the results from the workshop (Section 4), objectives and outlook of future actions (Sections 5 and 6) and some concluding remarks (Section 7).

## 2 The objective of the workshop

The key objective of this workshop was to bring together experts and stakeholders to define what the training and education requirements are from the different domains/target audiences, to identify state of the art as regards HPC training and education, and to examine how these different requirements may be met by existing or new European and national programmes on HPC training. The main outcome from the workshop is this report that provides guidelines for future education/training programmes at both national and European levels.

## 3 Identification of key stakeholders

Given the objective of the workshop it was recognised that representatives of the following communities and stakeholders had to be sought to achieve the desired outcome:

- HPC user communities;
- HPC application developers;
- HPC training providers (academia and industry);
- Providers of education in HPC (e.g. undergraduate and MSc programmes at universities).

An analysis of the geographical coverage and the make-up of the workshop participants is provided in the Annex 8.1.

## 4 Summary of results of the workshop

The agenda of the workshop was divided into three themes (the full workshop agenda is provided in Annex 8.2):

- HPC users - the HPC user community is the largest target audience group with the highest potential for short-term growth. This group is critical to demonstrate long-tail impact and wide adoption of HPC.
- HPC developers and user support - This community, although smaller in number compared to HPC users, is a critical component for adoption and effective hardware utilisation. They develop and maintain community codes and have the potential to quickly expand the user pool of novel resources, or support HPC users to maximise efficiency and productivity.
- HPC-related career paths - The full extent of HPC-related career paths is difficult to define, but there are clear indications that there will be difficulties in providing qualified people to satisfy the industrial activity related to this area. In order to support the acceleration of technological progress, there is a need to fortify and expand the structures that can generate pivotal, technology-enabling roles in the HPC space.

The first two themes focus on the needs and requirements of different communities represented by HPC users, and HPC application developers that also include HPC support staff. The third theme deals with both HPC education and support for HPC-related career path development. The agenda also includes preconceived questions for each theme that were devised to stimulate discussions among participants of the workshop; it was noted that not necessarily all of these questions would be raised and dealt with exhaustively given the time constraints of the workshop.

For each thematic session, experts and community representatives were identified and invited to participate as panellists that resulted in three panellists per session. The panellists were each given short periods of time at the beginning of each session to express individual views or points for discussion; they also had first priority in responding to questions from the agenda posed by the session chair. However, the majority of time of each theme was open to all participants of the workshop for interactive discussions.

### 4.1 Identification of needs and problems

With respect to each theme, there was an initial discussion to provide context for the needs of the various communities relevant to that theme. In the subsections 4.1.1 to 4.1.4 below, we outline the discussion as raised during the workshop.

#### 4.1.1 HPC users

- Support for HPC users, including training, requires local capabilities and sustained continuous support/training programmes. This is especially important for some communities who may have reluctance to embrace HPC. For smaller countries building this capability can take a lot of time and resources.
- There is wide recognition that PhD students who require and use HPC, but do not perform research on HPC, tend to be poorly trained in HPC. They are mostly trained solely on the job. HPC is not integrated in many conventional undergraduate courses, e.g. mechanical engineering, chemistry. We need a holistic approach to get students to use HPC naturally when engaging with scientific/engineering applications.
- We need “translators” who can talk to scientists but also understand HPC.



- In industry, the training needs of very large companies are not that much different to academia. But SMEs generally lack the expertise and their needs are more bespoke – special tools, interfaces, environments, etc.
- Industry training needs often demand short and tight time scales.
- High Performance Data Analysis (HPDA) will in a few years require more compute resources than HPC and many companies do not have the skills to address this.
- There is a need to teach HPC tools in universities from an early stage so graduates would use the same tools throughout their careers, e.g. from academia through to SMEs, industry.

#### **4.1.2 HPC application developers and support staff**

- Need to develop education and training for cohorts of developers who have knowledge of both hardware and software.
- High-end developers need to participate in refresher courses to keep up to date with technological advances. Technology providers have an important role to play in the provision of training and education.
- There is a need to ensure that training materials are kept up to date and incorporate modern standards and methods, otherwise it leads to poor HPC application design.
- The lack of continuous professional development and formal career recognition for high-quality trainers results in a high turnover of good trainers, ultimately leading to either shortages or inconsistent quality in the training provided.
- There is a need to develop training aimed at system administrators. There is a lack of qualified candidates in Europe and sites often have to hire externally. Workshops on this topic has proven popular in the past, e.g. in the Eastern Mediterranean as part of the LinkSCEEM project. Small universities and departments also demand such expertise to manage local systems. There is also proven demand from industry.

#### **4.1.3 HPC education**

- There are needs for MSc programmes as opposed to the 1-2 week accelerated approach (e.g. schools and workshops) that are generally more widely available to students, catering for those who prefer a “deep dive” into HPC. There is also scope for exploring new models, e.g. part-time, shorter programmes rather than a full-term MSc.

#### **4.1.4 Supporting HPC-related career paths**

- Support for internship and traineeship programmes (e.g. short-term positions at academic groups, HPC centres, technology vendors, ISVs that are open to European participants) are needed to facilitate mentoring and access to career paths across the EU.
- There is a need to develop training that takes into account all stages of people’s careers, i.e. lifelong learning, since HPC evolves at a fast pace and hence requires continuous re-training.

## 4.2 Lessons learned

Similar to Section 4.1, the workshop included a discussion of some of the existing activities and programmes across all three themes and the lessons learnt. The subsections 4.2.1 to 4.2.4 below provide a summary of the discussions.

### 4.2.1 HPC Users

- The landscape of support and training for HPC users is quite heterogeneous across European countries. There is a definite need for local interaction and personal contacts at the national level, in particular for new users, complemented by pan-European activities. The Competence Centres are a good initiative in this regard to provide visibility of the European landscape.
- Care should be taken when promoting technologies and ideas to new communities, in particular industry, in order to avoid too much hype and then failing to meet expectations – tailored messages may be required for different communities.
- The democratisation of HPC is happening – the cost of computing is getting lower for SMEs, but the problems appear to be cultural and structural. The Fortissimo approach is a good example of teaching SMEs to use numerical simulation and data analytics by first designing the solutions for them and then teaching others.
- In the field of biology, there have been some successful examples of integrating HPC teaching into the undergraduate/graduate curriculum, e.g. UCL, 200 biology students being taught currently and scalable up to 500 students. The integration needs to be done carefully – fully connected to the science and experiments.
- It is important to train users so they could work effectively on their own computer at home, which may involve dedicated on-site user support, which is supported by some CoEs, e.g. software expert visiting a research institute as part of the IS-ENES project funded by the EC.
- There are more and more online classes and learning opportunities opening up in the field, this addresses some of the scalability issues with face-to-face training, which cannot be easily scaled by nature. But there is consensus that a mix-and-match approach, i.e. both face-to-face and online offerings, is required. Especially for smaller countries, it is important to establish a national HPC support ecosystem that includes local training, which could then be supplemented by pan-European offerings.
- Online teaching should enable self-paced learning as well as opportunities for interaction among participants, and Q&A sessions, e.g. SPOCs (short private online courses).
- There is evidence from PRACE that local face-to-face courses are always in demand. However, low numbers travel to courses abroad, and cannot be substituted by online offerings. Trainer mobility, where trainers travel to provide courses, can address this issue at times but local expertise is still always needed.
- Face-to-face training is how some countries stimulate interest for HPC, and develop engagement with users, keeping them motivated.

#### 4.2.2 HPC application developers and support staff

- Case study: three major German HPC centres – High Performance Computing Center Stuttgart (HLRS), Juelich Supercomputing Center (JSC) and Leibniz Supercomputing Center (LRZ) – delivered 99 courses in Germany in 2018:
  - Courses were mostly focused on application/software developers in academia.
  - 19 out of 99 courses were conducted through PRACE.
  - Programme aims to bring participants from zero HPC knowledge to 100%, i.e. multi-teraflop, petaflop, exascale (typically at least petaflop). The basic principles of the programme are largely the same over time, with the concept of “zero to 100%” based on somebody who already has full knowledge of the domain (e.g. the physics), numerical methods.
  - 5-10% participants from industry; an expert opinion is that industry prefers to hire those who are already trained from academia.
  - Developed a train-the-trainers programme since 2013 for state of the art OpenMP and MPI.
  - Trained 41 trainers with 7 potential long-term success stories (in Belgium and Austria). For the others, some lack the organisational background to provide HPC training; trainers are “lost” over time for a variety of reasons, e.g. fixed term contracts.
  - Career progression for trainers is difficult under typical academic settings; no research profile for path to professorship. Although this may not always be the case for every university.
- There is capacity for technology providers to contribute to HPC training, e.g. Intel contributes to some MSc programmes in HPC. Vendors sometimes contribute to HPC short courses organised by academia, e.g. extreme scaling workshops.
- The CERN School of Computing supports 40-50 undergraduates for 2-3 months placements. Acceptance into the programme is highly competitive and students are highly motivated. This type of initiative has the potential to get the community excited.
- Summer Schools would be a suitable model for professional education, e.g. industry participation of 33% for data science.
- There is the VI-HPS project that provide entry-level courses on performance engineering using tools developed by its partners. Courses are distributed over Europe and the U.S., many supported by PRACE.
- The POP CoE organises more in-depth training (in collaboration with EoCoE, E-CAM, MaX) focusing on in-depth analysis of application performance, working with small groups of developers on own codes.

#### 4.2.3 HPC education

- Case study (Italy): MSc programme has been running since 2013 for approx. 15 students per year. Majority are PhD students with strong competencies in their domains. The programme is jointly funded from industry and many eventually go into industry. There is a separate programme to bring undergraduates to the required level.
- Case study (Ireland): MSc programme for more than 20 years, initially to grow the HPC community and then coincided with the arrival of the first large HPC systems in Ireland.

Requires good undergraduates, and there has been consistent interest from the financial sector. The final projects are a very important component.

- Most MSc graduates go into data science positions, 25% take up PhDs, a good number go into finance industry.
- Case study (UK): Well established programme that takes in around 60 students per year, with about 50% of students from outside Europe.
  - The university will run fully online HPC MSc programmes from 2020.
- HPC MSc students typically enrol to pick up new technical skills that they can apply to their primary domain, they are not usually interested in following a solely HPC career path.
- There is a market for more MSc programmes, potentially exploring different paths (infrastructure, application development, etc.). But coordination at the European level would be useful to discuss offerings. Providing students opportunities to travel would be of use also.
- Part-time MSc programmes would be useful (for on-the-job training) and/or shorter programmes, e.g. diploma or certification. But resources for delivery (e.g. where evening teaching is required) could be problematic for some.
- Companies are generally supportive of MSc programmes without dictating what they need in terms of the content. Some programmes are, however, designed with a lot of industry influence.
- The MSc programmes in Ireland, Italy and the UK can attract anywhere from 30-50% of students from outside Europe.
- There may be difficulties in pushing for European level Masters HPC programmes as this impacts on local and national legislation. Integration into existing university programmes is much more feasible.

#### **4.2.4 Supporting HPC-related career paths**

- Internships are a good way to help smooth career transitions.
- Well defined career paths, and in particular career progression, matters more for those who have already embarked on a career, as compared to those who are considering taking an HPC path. There are many cases where people organically, rather than deliberately, end up in HPC roles.
- While HPC centres need to provide clearer career paths, visibility of career paths should encompass more than just HPC centres, e.g. what are the professional profiles out there? For example, long term application development.
- Internships play an important role and of genuine interest to companies large and small, e.g. banks, insurance companies, SMEs. However, it takes a lot of effort to build up contacts and engage with companies to offer internship opportunities for large groups of students.
- Support for academic programmes along with mobility programmes (e.g. HPC-Europa, Summer of HPC, teaching staff) would help to expand access to career paths across the EU.

## 5 Objectives of future actions

### Support for existing stakeholders and initiatives

- Support the existing CoE training programmes that offer a rich mix of activities targeted at different levels and domains of HPC users; from short courses, hackathons, online courses to mobility programmes. Coordination at a European level for this type of activities is needed.
- Many different European projects need to do training, including PRACE via its training centres, schools, workshops and MOOCs. EuroHPC should recognise that training is critical and must be sustained.
  - There are already organisations and institutions such as PRACE, CoEs, national centres, universities that are well placed as vehicles for training. Competence Centres will join this list in 2020.
- Technology providers should be involved in training provision for both academia and industry to showcase the latest technology capabilities and best practices.

### HPC education

- More bridges/interaction are needed between education providers and HPC experts.
- Develop tracks in universities to produce the needed cohort of people, from HPC users on one end of the spectrum, the “translators” who occupy the middle ground and hard-core gurus at the other end. A healthy numerical balance on this spectrum needs to be maintained.
- Integrate HPC into undergraduate and graduate curricula, by first developing some good examples to attract attention, followed by approaching the right people in the university system. This may be more difficult for undergraduate curricula that have certain inertia to changes.

### HPC users and developers

- Efforts should target the whole of the academic spectrum including social sciences, digital cultural heritage, etc.
- Develop dissemination activities targeting domain-specific conferences and symposia to enhance awareness of HPC, which is still lacking for user communities.
- Make sure training and support is provided not only for HPC, but also on HPDA that will require more and more resources and is particularly relevant to industry.
- Ensure that training does not just focus on the current level of competency of the trainees. Trainees should be exposed to more advanced subjects, so they know the potential benefits and rewards.
- Provide support to train developers and support staff in software engineering tools and techniques (e.g. version control, continuous integration), even though this may be seen as peripheral to HPC. It's important that HPC developers and support engineers are skilled in proper scientific software development and general software development best practices.
- Develop dedicated people to engage with industry to create contacts, build awareness and examine evolving, bespoke training requirements.
- There is insufficient appreciation of the role of application developers in research, particularly when it comes to exascale. There is scope for a program similar to the

SciDAC<sup>1</sup> initiative in the US that can give prestige to this critical role and make it an attractive career path.

### **Trainers**

- Maintain an ongoing pool of HPC trainers by improving continuous professional development.
- Ensure that training material can be re-used and kept up to date.

### **Mobility programmes, traineeships, Masters programmes**

- Support existing and new mobility programmes for training and mentoring, e.g. the HPC-Europa programme to fund HPC-related research visits, the PRACE training mobility fund to support trainers teaching abroad, software expert visits to research institutes.
- Support internship programmes such as the PRACE Summer of HPC programme that places students at HPC centres around Europe. Other types of programmes that cater for placements in industry from academia, or *vice versa*, should be considered.
- Develop traineeship opportunities for future system administrators, e.g. apprentices are placed in production HPC centres for extended periods.
- Certification would be useful for on-the-job training. Technology providers would see this as a positive attribute.
- Explore the development of MSc programmes that may incorporate traineeship opportunities and that target different tracks (e.g. application development, infrastructure), with potentially part-time and/or shorter programmes (e.g. diploma, certificates).

## **6 Outlook on future actions**

- Aim for tighter integration of domain science with numerical and computational methods for education/training and research – European coordination, rather than national ones, would be more useful.
- Consider means to support career progression and opportunities for HPC trainers, in particular those based in academic institutions.
- There are perhaps more people providing HPC support than those providing training. Consider ways to encourage more participation from support staff to participate in, or move into, training activities.
- For the longer term, the quality assurance of training being provided needs to be properly addressed.
- More work is needed to elucidate diverse career paths and professional profiles (across academia, industry, HPC centres), along with their expected skills requirements, that could be addressed in future at a European level.
- Develop a framework for European level communication and coordination of HPC training activities – regular ecosystem meetings, increased visibility of activities, discussion on methodologies and what works (or not).

---

<sup>1</sup> <https://www.scidac.gov/about.html>

## 7 Conclusion

With respect to user communities, it is clear that CoEs, in collaboration with PRACE and HPC centres are providing a wide spectrum of activities. The level of scalability of this is questionable though, as at present there is difficulty in maintaining qualified training personnel even in countries where HPC is well established.

For countries where HPC is not so established, such as some EU13 countries, sufficiently qualified personnel would have to be imported and ideally go on to train local personnel. If the market is not there to provide qualified personnel, then the quality of training provided suffers as a result. There would appear to be a clear need to provide for a program that can deliver a consistent level of user training for all EuroHPC countries, which includes building consistent local competence all around. Where appropriate, local initiatives can refer people to European initiatives (such as PRACE and the CoEs) for advanced training. For system administration skills there is also a lack of training opportunities. This is another activity that could be coordinated at a European level and include extensive internship opportunities.

As regards the development of scientific software for HPC, while training opportunities exist there is typically insufficient appreciation in academia for the role this plays in modern research. In many cases, engaging deeply with application development comes at the cost of publications, and ultimately a sustainable career. A programme similar to the SciDAC program in the US could provide sufficient prestige for researchers in this role to allow them to more effectively compete for permanent research positions.

Many of the workshop participants were not fully aware of the extent of training initiatives across Europe, particularly when it comes to formal HPC education. The feedback from the participants for the overall event was that they would indeed like to hold a similar event again in the near future. One key point raised, for example, was that the training required that relates to technology providers was not really discussed in more detail.

In general, with the experience gained from the first iteration of this workshop, it is clear that there is significant scope for increasing the dialogue among HPC training and education providers. The level of experience and expertise among the community is considerable, and the discussion was open with respect to the challenges and frustrations that are being faced. In an era where we expect the training and education requirements for HPC to expand rapidly, the communication among, and recommendations from, such a group has the potential to provide substantial impact.

## 8 Annex

### 8.1 Participants

The workshop hosted 31 participants (excluding one member of staff from the EC) from 14 countries based on institution. The geographic distribution is shown in Figure 1 below. This included 5 participants from 5 different EU13 countries.

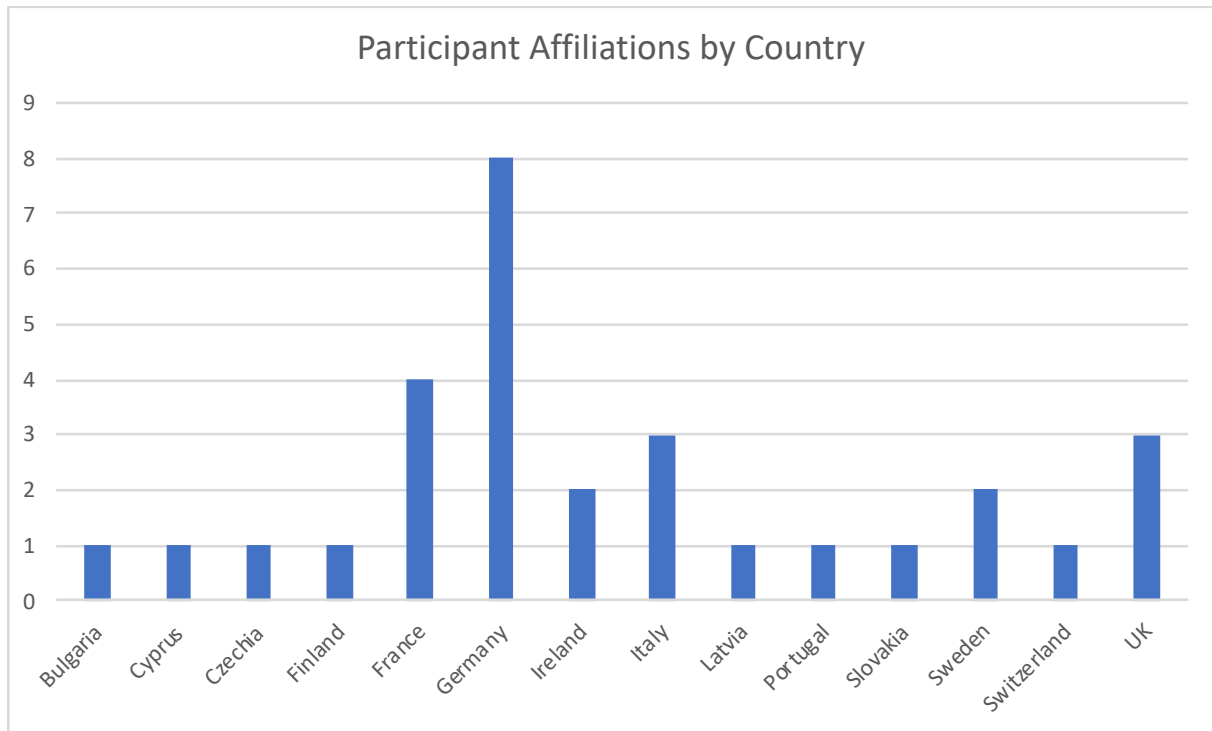


Figure 1. Geographic distribution of participants by country of institution

The participants were invited to the workshop to represent the interests of various groups (user communities, application developers, training and education providers) in academia and industry. They came from a diverse range of organisations and projects:

- 10 participants from CoEs
- 5 participants from EU13 countries (mainly local/national HPC service providers)
- 4 participants from industry and industry-related interest groups (e.g. Teratec, ETP4HPC)
- 3 participants from FocusCoE
- 3 participants from PRACE
- 3 participants representing HPC education providers, from institutions (e.g. universities) running HPC-focused MSc programmes
- 2 participants from national HPC centres and training providers
- 1 participant from the BDVA project
- 1 participant from the EC

It is important to note that some of the participants are involved in multiple projects (e.g. both CoE and PRACE) as well as projects not listed above (e.g. both PRACE/CoE and EXDCI-2), hence the collective knowledge of participants is more diverse than what is indicated.



## 8.2 Workshop agenda

### European HPC Training Stakeholder Workshop

8<sup>th</sup> October 2019

European Commission premises in Room 0/54, Avenue de Beaulieu 33, Brussels.

### Workshop Background

HPC training plays an important role in sustaining a steady stream of researchers and developers who can exploit the full potential of supercomputers. Its role is set to grow in the coming years as Europe looks to develop the next generation of technologies, applications and systems towards exascale, as well as enhancing uptake of HPC from industry and achieving excellence in HPC application delivery and use. This workshop is aimed at gathering experts and stakeholders both to define what are the training requirements from different domains/target audiences, and to examine how these requirements may be met by existing or new European and national programmes on HPC training.

### Tentative Timetable

09:00 - 09:30: Introduction from EC - overview of participants - objectives for the day

09:30 - 10:30: Theme 1: Developing the user community

10:30 - 11:00: Break

11:00 - 12:00: Theme 1: Developing the user community

12:00 - 13:00: Lunch

13:00 - 14:30: Theme 2: Developing the HPC developer and user support community

14:30 - 15:00: Break

15:00 - 16:30: Theme 3: Supporting HPC-related career paths

16:30 - 16:45: Closing remarks

*Modus operandi* for each theme:

- 3x panel members (opening statements & discuss questions from chair)
- 1 chair (raise pre-defined questions from samples below, moderate discussions, questions/comments from floor, distil proposals / recommendations)
- FocusCoE personnel will keep minutes of discussions, in particular try to highlight any common/agreed positions
- Proposed structure of the workshop report is on the last page; discussions during the workshop will mainly feed into Sections 3-5.

# Theme 1: Developing the HPC user community

## Motivation

The HPC user community is the largest target audience group with the highest potential for short-term growth. This group is critical to demonstrate long-tail impact and wide adoption of HPC.

## Relevant questions

The needs and requirements of the community

- Which HPC users should be targeted? All the way from novice to advanced users? Which user communities?
- Are industry needs different to that of academic HPC users?
- What is the mapping between commercial software environments and EU extreme computing resources?
- Do we need to permeate HPC awareness and education into university undergraduate or graduate curricula? If so, how to do this for numerous universities?

Current landscape and challenges

- What are the current activities? CoEs, PRACE, universities (MSc, PhD), commercial training...what else? What works and what doesn't?
- How do we address scalability of training? e.g. increase availability of certain courses (which ones?) across Europe by ramping up the numbers of trainers.
- Should we aim for equitable access to training (courses, material) across Europe? Should EU fund mobility for those HPC users who need travel support?
- How do we ensure equitable access to training infrastructure across Europe? Who is responsible for maintaining such infrastructure? How do we lower the technical barriers of moving from one HPC site to another?
- How do we ensure trainees have equitable access to computational resources? Do we need to minimize diversity in training environments?

Looking forward...

- How do we reconcile European vs national training efforts/initiatives/activities?
- How do we estimate the initial effort required, track progress and develop/measure success? i.e. quantity and quality of courses? Number of people trained? More resource-intensive impact assessments (longitudinal surveys, focus groups, interviews)? Publications?
- Should we invest more into remote learning initiatives? Which type (e.g. MOOCs, SPOCs, hybrid online+F2F courses)? What makes a successful remote learning experience? What pitfalls do we need to avoid? What platforms are needed to support remote learning?
- Who pays for the expenses related to accessing training? Who pays for the development, maintenance and delivery of courses?
- Would a "HPC Passport" type initiative be of benefit to industry?

## Panelists

1. Constantia Alexandrou (The Cyprus Institute)
2. Daniel Verwaerde (Teratec)
3. Sophie Valcke (Cerfacs)

## Theme 2: Developing the HPC developer and user support community

### Motivation

This community, although smaller in number compared to HPC users, is a critical component for adoption and effective hardware utilisation. They develop and maintain community codes, and have the potential to quickly expand the user pool of novel resources, or support HPC users to maximise efficiency and productivity.

### Relevant questions

#### The needs of the community

- Which type of developers should be targeted? Novice to advanced developers? Across all disciplines? Do we want to foster a larger new pool of developers (to what extent?) or simply upskill existing ones?
- Do we need different formats of training for different levels of developers? e.g. traditional F2F course for beginners, hackathons or highly-focused/challenge-based training for more advanced developers?
- Do we need more activities focused on system administrators, HPC hardware specialists (e.g. those who'd develop new HPC technologies, R&D staff at vendors)?
- What are the industry expectations?

#### Current landscape and challenges

- What are the current activities? CoEs, PRACE, universities (MSc, PhD), HPC Competence Centres, commercial training...what else? What works and what doesn't?
- Do we need to ensure much wider awareness and adoption of software development best practices? Are the relevant tools available across European HPC resources to support this?
- Should we aim for equitable access to training (courses, material) across Europe? Should EU fund mobility for those HPC developers who need travel support?

#### Looking forward....

- How do we establish and support Research Software Engineer as a viable career choice across Europe?
- Who pays for the expenses related to accessing training? Who pays for the development, maintenance and delivery of courses?
- Should software infrastructures be included in this discussion (e.g., software installation frameworks such as EasyBuild and Spack; CI infrastructures)?

### Panelists

1. Brian Wylie (Juelich Supercomputing Centre)
2. Marie-Christine Sawley (Intel)
3. Rolf Rabenseifner (HLRS)

## Theme 3: Supporting HPC-related Career Paths

### Motivation

The full extent of HPC-related career paths are difficult to define, but there are clear indications<sup>1</sup> that there will be difficulties in providing qualified people to satisfy the industrial activity related to this area.

In order to support the acceleration of technological progress, there is a need to fortify and expand the structures that can generate pivotal, technology-enabling roles in the HPC space. For example, in academia there is an increasing appreciation for the role of *Research Software Engineer*, who combine professional software expertise with an understanding of research.

### Relevant questions

Career paths, workforce demands

- What career paths are well defined? Are there opportunities to formalise others?
- How do we expand access to these career paths across the EU (while avoiding brain-drain)?
- What about gender balance?
- How can we identify and address (future) industrial demand in a proactive, rather than reactive, way?
- What about the systems side of HPC? i.e. the HPC technology industry, e.g. vendors, hardware specialists.
- How does the EuroHPC ecosystem feature in this goal?

Instruments: programmes such as MSc, traineeships, mentorships, mobility

- What are the challenges in growing the formal education component (M.Sc. programs, etc.)? Is there demand? If not, how do we grow it?
- Is professional certification a possibility? What are the challenges/benefits?
- Would mentorship/traineeship mobility programmes help? e.g. staff/student visiting/hosting programmes, HPC Europa?

### Potential Panelists

1. Mike Peardon (Trinity College Dublin)
2. Stefano Cozzini (SISSA & ICTP, eXactLab co-founder)
3. Weronica Fillinger (EPCC)

---

<sup>1</sup> See JRC Technical Report "Academic offer and demand for advanced profiles in the EU *Artificial Intelligence, High Performance Computing and Cybersecurity*"