



PRACE introduction

EUDAT – PRACE Summer School on managing scientific data from analysis to long term archiving, 23-27 September 2019, Trieste, Italy

Simon Wong, Irish Centre for High-End Computing (ICHEC), Ireland

Leon Kos, University of Ljubljana, Slovenia

Overview

What is High Performance Computing (HPC)?

Overview of PRACE

What you'll learn

- PRACE services (HPC access, support, training)
- Getting around the HPC environment
- Tools and services that are linked to EUDAT

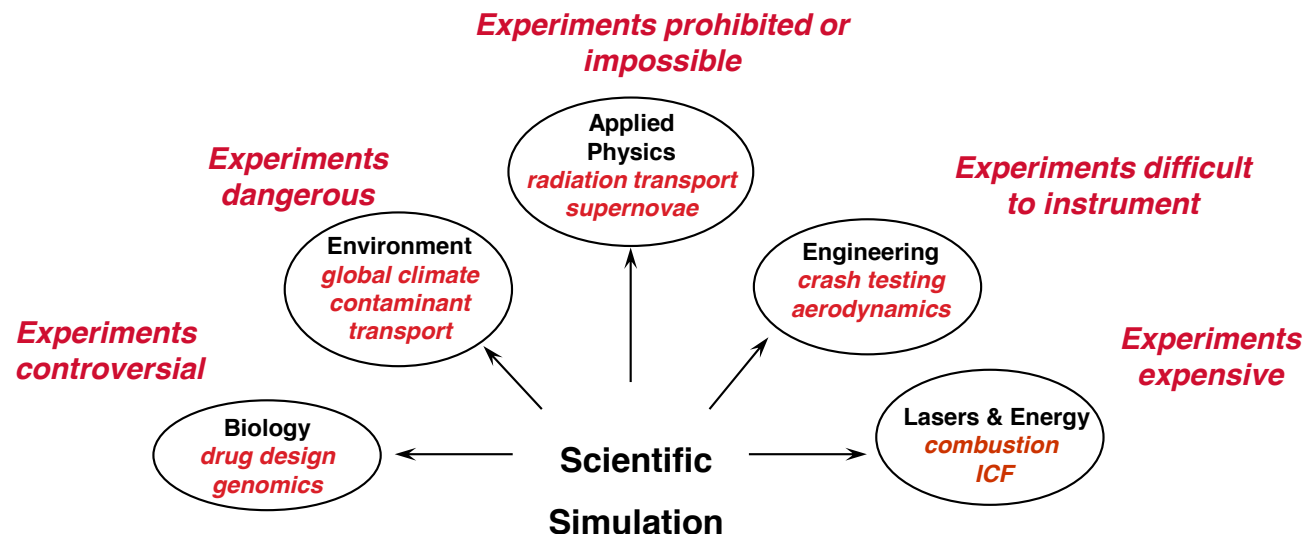


High Performance Computing / Supercomputing

HPC – aggregation of computing power to achieve much higher performance than a single computer, and orchestrating software on this system to solve large problems.

Most scientists look to HPC...

- My compute workload is taking too long on my laptop!
- Here, use this machine to do your experiments/analyses.



Typical route to access systems managed by research institute, university, national HPC centre, etc.




Apply for resources (via those managing the system, or sometimes via funding agencies).



Review for technical feasibility and scientific merit.



Receive an allocation of HPC resources (e.g. core hours).

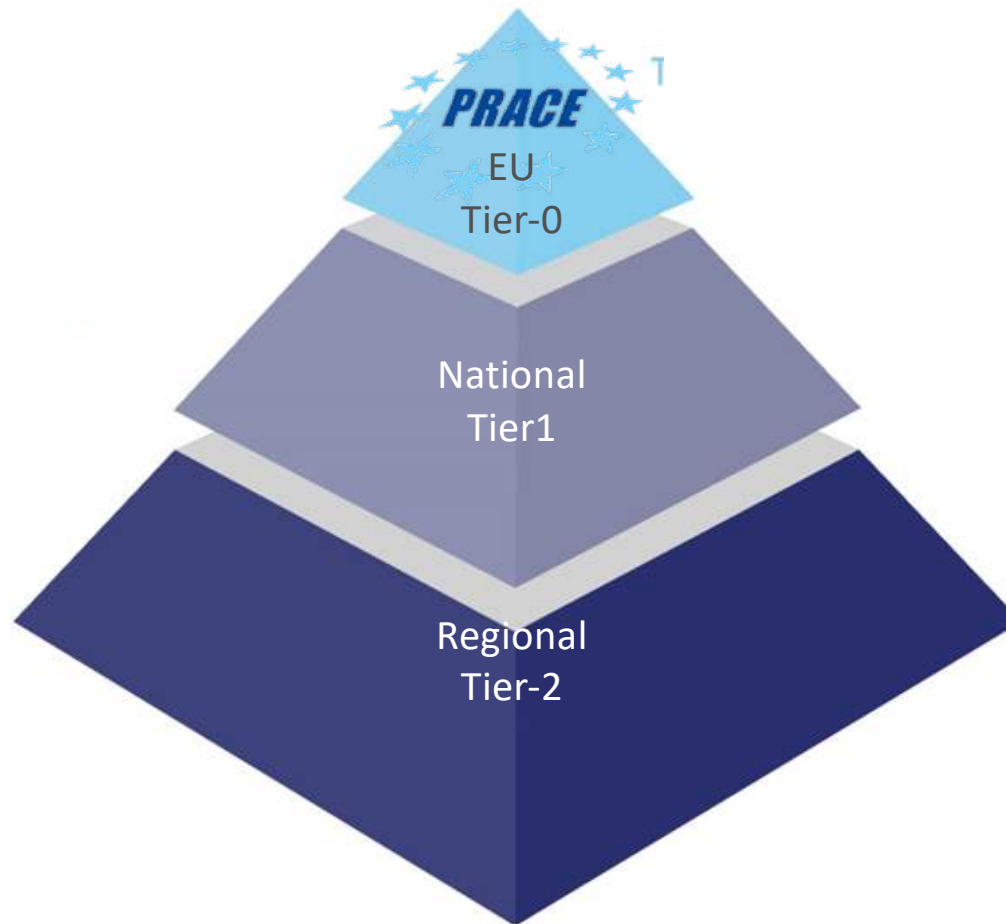


Log into machine; set up data, software applications
→ submit compute jobs.

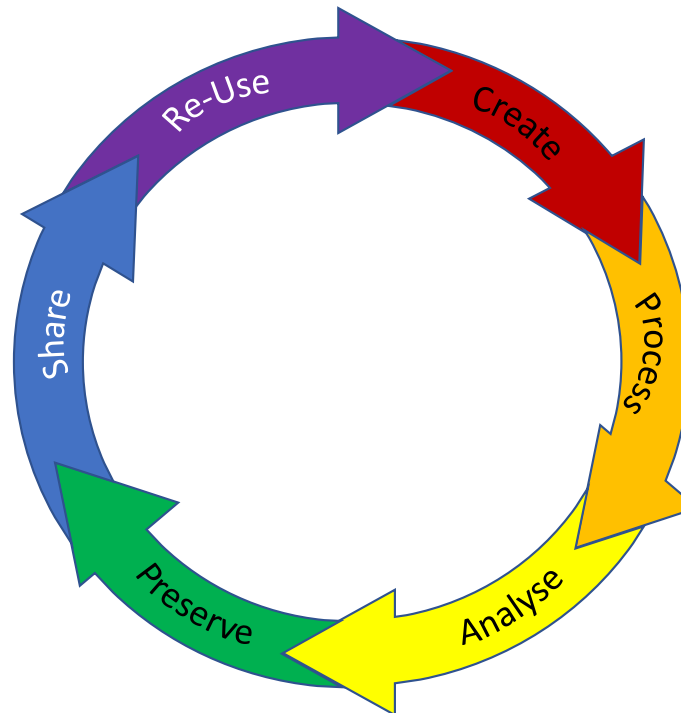



(Take key data out), analyse/summarise output and publish.

What if I require LARGE amounts of resources?



EUDAT, PRACE – Data Life Cycle



- Access to infrastructures
- Support for simulations and processing workloads
- Visualisation
- Training

Partnership for Advanced Computing in Europe (PRACE)

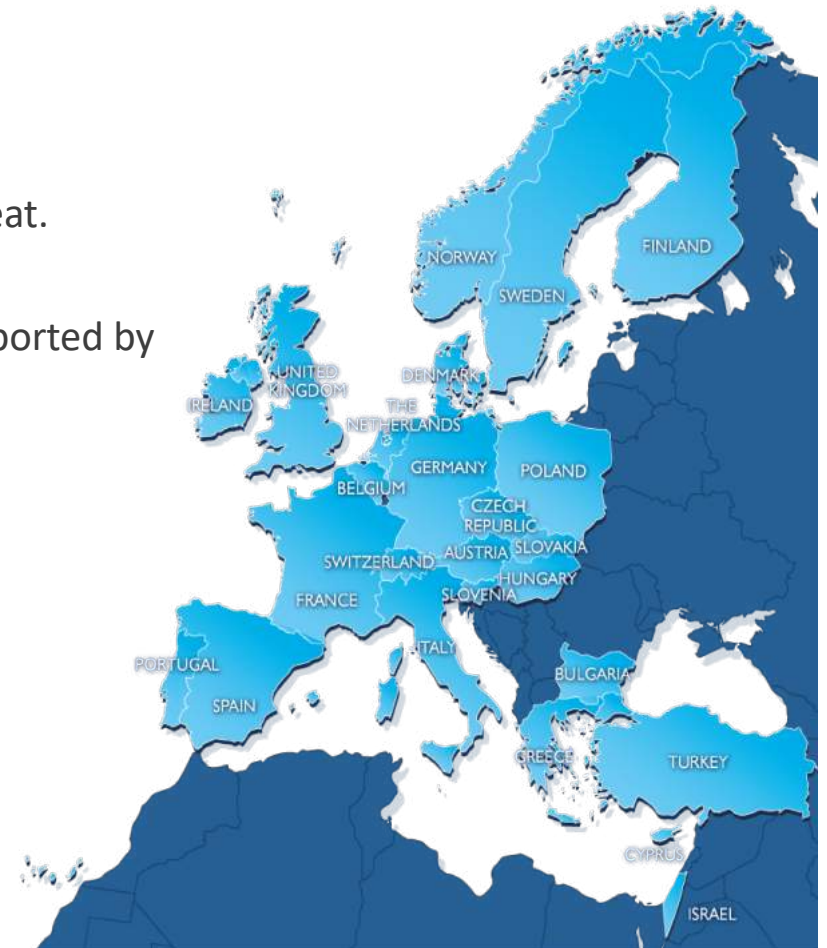
PRACE is an international not for profit association under Belgian law.

26 members, 2 observers.

Governed by the PRACE Council in which each member has a seat.

Funded by members as well as through a series of projects supported by the European Commission.

Computing resources are made available by a group of members – hosting members.



PRACE members

Hosting Members

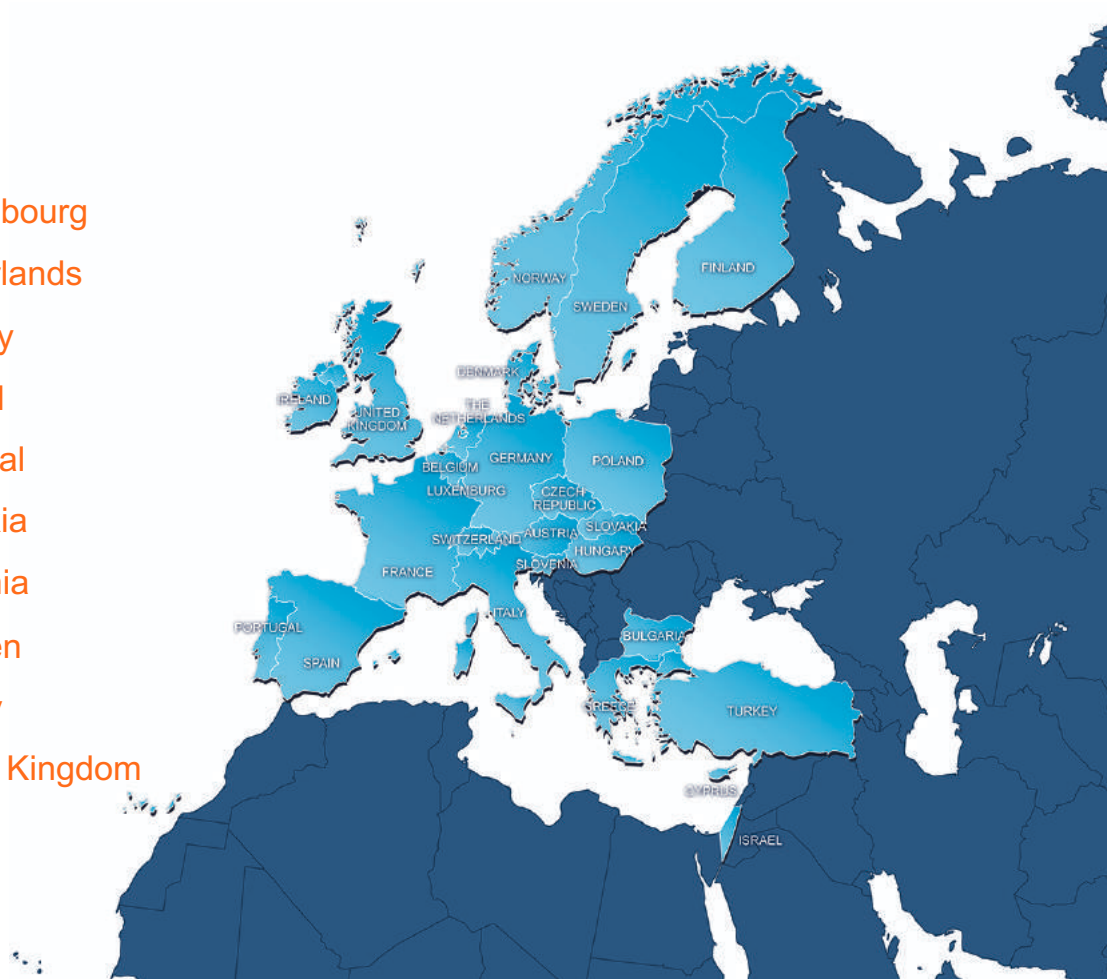
- ▶ France
- ▶ Germany
- ▶ Italy
- ▶ Spain
- ▶ Switzerland

General Partners (PRACE 2)

- ▶ Austria
- ▶ Belgium
- ▶ Bulgaria
- ▶ Cyprus
- ▶ Czech Republic
- ▶ Denmark
- ▶ Finland
- ▶ Greece
- ▶ Hungary
- ▶ Ireland
- ▶ Israel
- ▶ Luxembourg
- ▶ Netherlands
- ▶ Norway
- ▶ Poland
- ▶ Portugal
- ▶ Slovakia
- ▶ Slovenia
- ▶ Sweden
- ▶ Turkey
- ▶ United Kingdom

Observers

- ▶ Croatia
- ▶ Romania



PRACE Tier-0 systems (2019)



MareNostrum: IBM
BSC, Barcelona, Spain
#29 Top 500



Piz Daint: Cray XC50
CSCS, Lugano, Switzerland
#6 Top 500



NEW ENTRY 2018/2019
SuperMUC NG : Lenovo
cluster GAUSS @ LRZ,
Garching, Germany #9
Top 500

NEW ENTRY 2018
JUWELS (Module 1):
Bull Sequana
GAUSS @ FZJ, Jülich,
Germany #30 Top 500



NEW ENTRY 2018
JOLIOT CURIE : Bull Sequana
GENCI @ CEA, Bruyères-le-Châtel,
France #47 Top 500






MARCONI: Lenovo
CINECA, Bologna, Italy
#21 Top 500

Hazel Hen: Cray
GAUSS @ HLRS,
Stuttgart, Germany
#34 Top 500
NEW ENTRY soon:
HAWK



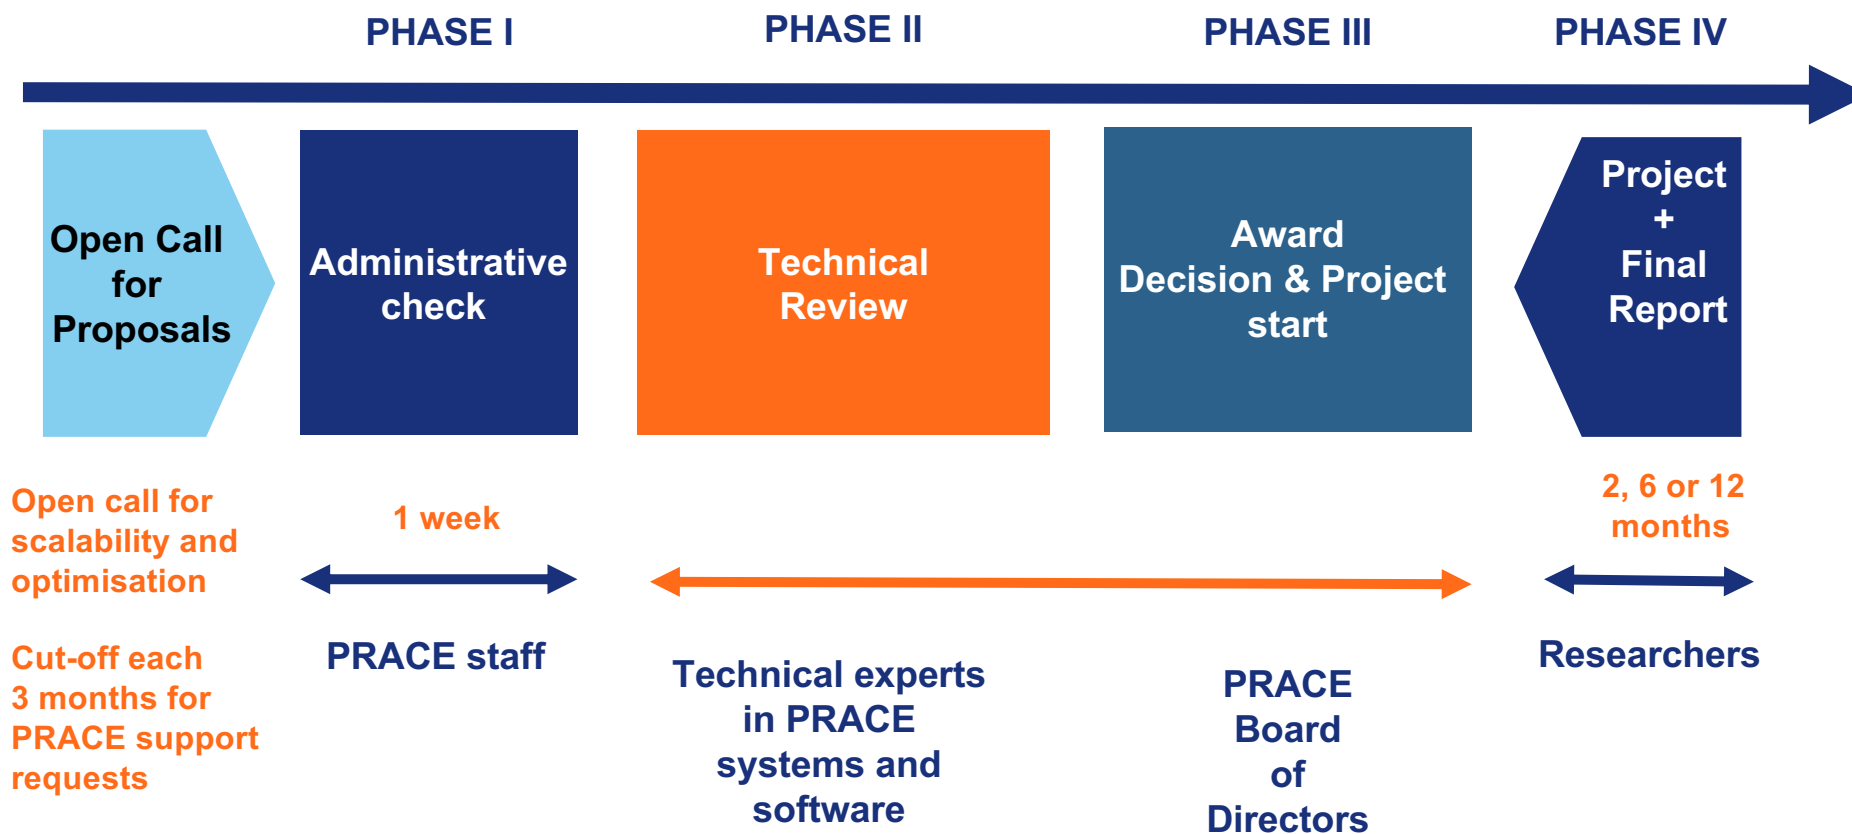
**Close to 110 Petaflops
total peak performance**

Access to PRACE resources

-  **Free-of-charge** required to **publish results at the end of the award period**
-  **Preparatory Access (2 to 6 months)**
-  **SHAPE Programme (2 to 6 months)**
-  **Project Access (12, 24 or 36 months)**
-  **Centers of Excellence : 0,5 % of the total resources available for the 11th call for CoE**

**Criterion:
Scientific Excellence
Assessed by an
improved review
process**

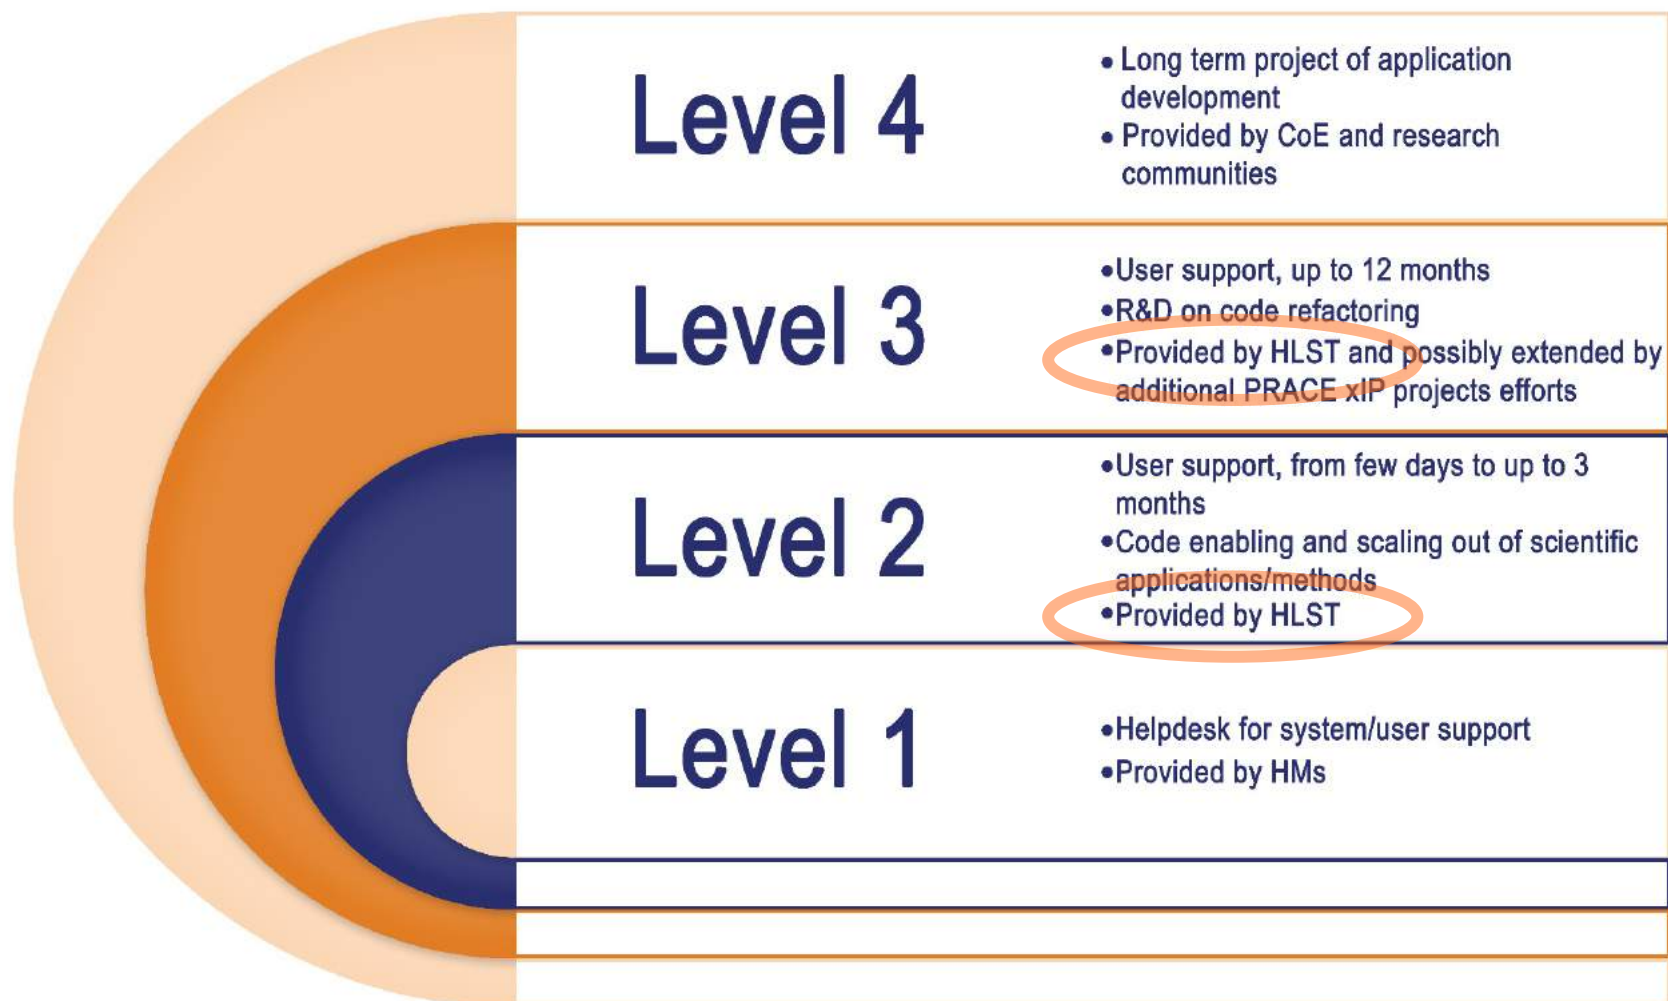
Preparatory access



Project access



Support for users and communities



Access to Tier-1 systems (DECI programme)

- Tier-1 systems are operated in PRACE sites distributed in ~16 different European countries.
- Tier-1 systems collectively provide a total of peak performance of 29 PFlop/s.
- Most of the Tier-1 systems are accelerated with GPU cards or Intel co-processors (Xeon Phi).



EUDAT and PRACE are joining forces to help research communities gain access to high quality managed e-Infrastructures whose resources can be connected together to enable cross-utilization use cases and make them accessible without any technical barrier.

Training

Around 100 F2F training events per year

- PRACE Training Centres
- Seasonal Schools (Int. HPC Summer School)
- Collaborative events

MOOCs

- Supercomputing (14 Oct 2019)
- Python in HPC (9 Sep 2019)
- Managing Big Data with R and Hadoop
- Defensive programming and debugging

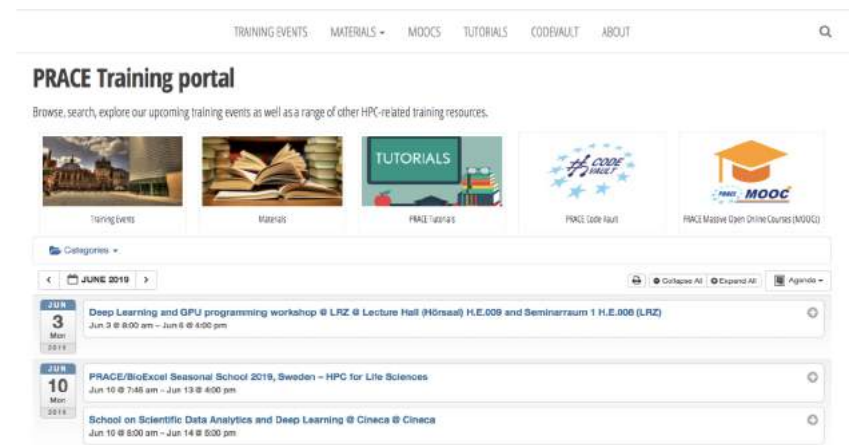
Summer of HPC


- Undergraduate to early graduate students

Training Portal

- Courses, materials, MOOCs, code samples


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





Other relevant activities

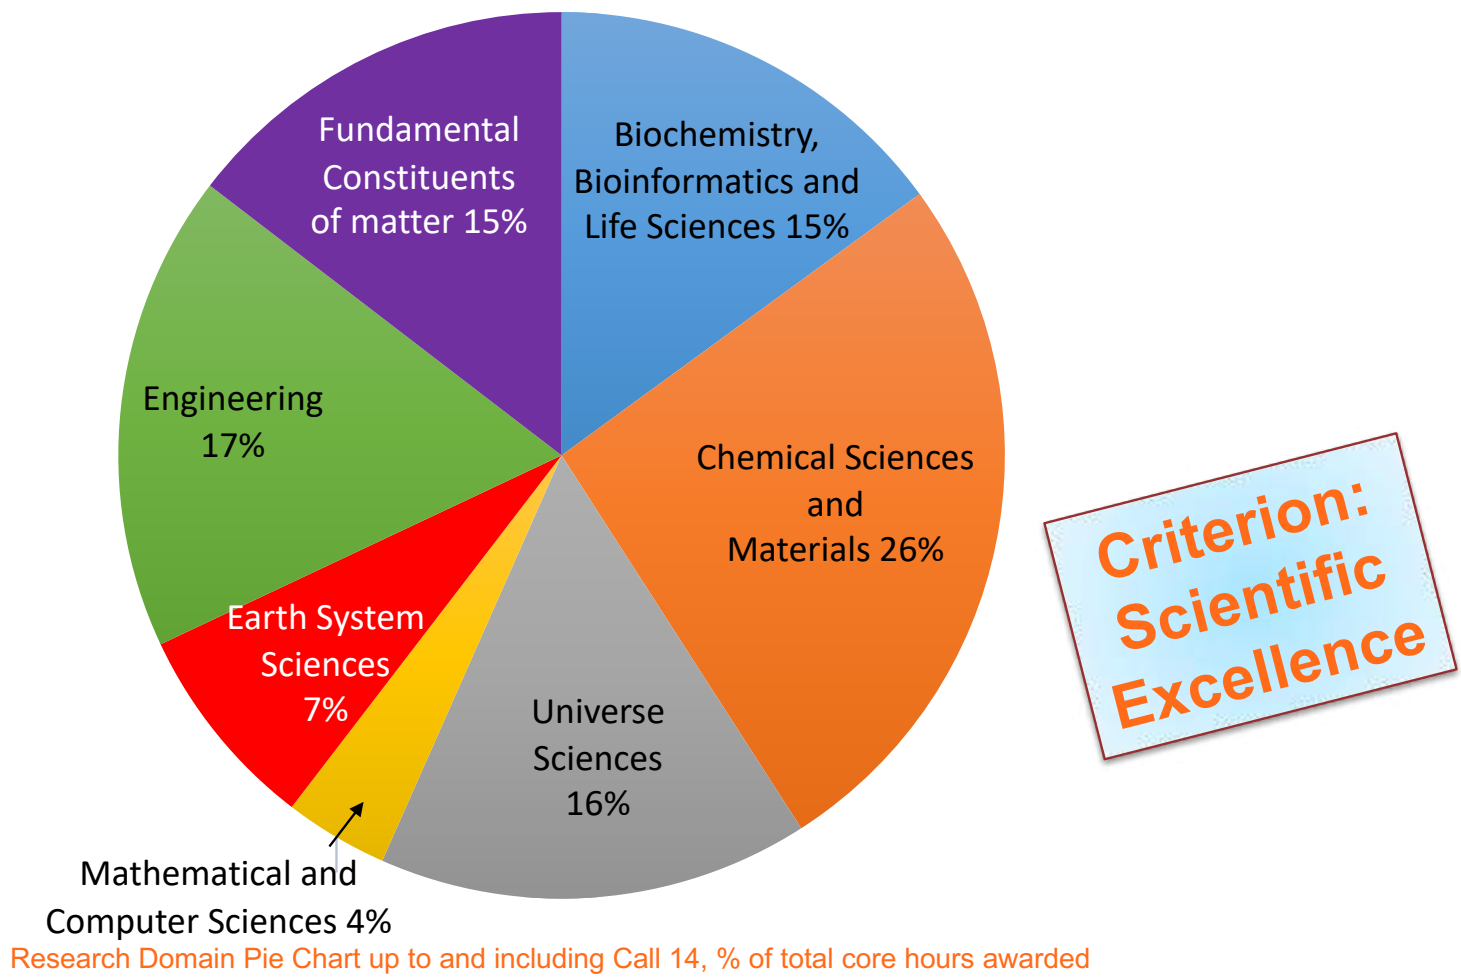
- Internal standardisation and harmonisation activities to ensure a level of coherence across PRACE systems.
- Code and application enabling activities – how to make existing codes more efficient, and scale up in order to be able to tackle larger problems in reasonable time.
- Investigate emerging HPC technologies for scientific applications.
- Develop next generation software solutions for extreme scaling.
- Provide HPC access and support to industry via the SHAPE programme.



Main achievements

- **688** scientific projects enabled
- **>21 000 000 000** (thousand million) core hours awarded since 2010
- Of which **63%** led by another PI nationality than the HM
- R&D access to industrial users with **>50** companies supported
- **>12 000** people trained through PRACE Training
- **~110** Petaflops of peak performance on **7** world-class systems
- **26** PRACE members, including **5** Hosting Members
(France, Germany, Italy, Spain and Switzerland)

Supporting many scientific domains



Scientific Case for future supercomputing

- **Scientific applications**
 - urgent need for more compute cycles, AND
 - huge demands memory bandwidth & I/O
- **Need new approaches**
 - scaling via ensembles, deep learning, and statistical models
 - systems able to handle tens of thousands of active jobs and large I/O requests
- **Software & algorithms** take longer to change than hardware
 - PRACE & Europe need a much more ambitious strategy to develop the SW part of next-generation computing



“Future Infrastructures and operations will need to be much more diverse to support HPC, Data Science and different types of accelerators - but we also need to avoid fragmentation.” - PRACE Scientific Case

Thank you. Questions?