

PRACE User Forum Open Discussion



EuroHPC
Joint Undertaking

**EuroHPC
Summit Week
2022**

From 22 to 24 March 2022 | Paris, France

#EHPCSW

Overview

- PRACE User Forum introduction
- Recap: user overview
- Open discussion I: general matters
- Open discussion II: the future of the PRACE User Forum

Role of the User Forum

- A communication channel between the user community and the resource provider
- A forum in which users can share experiences
- Programme committee
 - organizes user fora
 - Give users a voice in the European HPC debate
 - reports back to PRACE
 - ✓ identification of user problems
 - ✓ recommendations to meet user needs

User Forum Membership

- All researchers who have used PRACE resources within the past 5 years are considered to be members of the User Forum.
- The Program Committee of the User Forum is drawn from the User Forum membership and creates consensus within the User Forum regarding proposals to be brought forward to PRACE on behalf of the user's community.
- We happily consider new members for the Program Committee.

User Forum Committee


- Current members
 - Derek Groen (Chair Life Sciences, UK)
 - Carmen Domene (Vice-chair, Chemistry and biophysics, UK)
 - Troels Haugbølle (Chair, Astrophysics, Denmark)
 - Gabriel Staffelbach (Engineering & Energy, France)
 - Gustavo Yepes (Astrophysics, Spain)
 - Jorge Vieira (Plasma physics, Portugal)
 - Koen Hillewaert (Engineering & Energy, Belgium)
 - Marc Baaden (HPC, France)
 - Teresa Parra-Santos (Engineering, Spain)
 - Turlough Downes (Astrophysics, Ireland)
 - William Sellers (Biomechanics, UK)

European Tier-0 users


Based on survey of PRACE Tier-0 PIs carried out by the user-forum in 2019

Core hours used per year

- < 100,000
- 100,000 - 1,000,000
- 1,000,000 - 10,000,000
- 10,000,000 - 100,000,000
- >. 100,000,000




Architecture used for the codes



- 10 – 100M CPUh / yr typical usage
- 12% use +100M / yr
- In house codes, or "standard" field specific codes
- Increasing adaption of accelerators, but still minority of use cases.

Type of code running on PRACE



View of European Tier-0 users

Based on survey of PRACE Tier-0 PIs carried out by the user-forum in 2019

- PRACE in 2019 was the only credible option for Tier-0 level science. To quote from the survey

“In Europe there is no other alternative than PRACE if one needs allocation of size of order 100M core-hr”

- European-level resources crucial for the digital agenda
- PRACE is well respected for their fair and balanced application and review process
- A key to the success of PRACE is that the infrastructure has been developed bottom-up with input from scientists and computing centers

Open Discussion: General matters

- Review process and grant allocations
 - What will change during the transition to EuroHPC?
- Do we need common storage, data intensive and high-throughput computing to support:
analysis, reuse, and publishing of Tier-0 project data?
- Industrially oriented projects
 - open/closed data and reproducibility
 - scientific vs technological progress
- How can PRACE/EuroHPC best support applied scientists from academia and industry in their use of the infrastructure and in transitioning to accelerated architectures ?
- Other remarks / needs / problems ?

Open Discussion II: Future of the User Forum

- EuroHPC has now kicked off, and PRACE will transition away from mainly providing allocations towards mainly providing user support and training.
 - Although it will provide cycles to some extent, and may still do peer-reviews.
- EuroHPC is in the process of establishing a User Forum, which overlaps with the communities in the PRACE User Forum, but comprises a slightly different set of users.
 - We mainly represent research-oriented HPC users.
 - PRACE users are present globally.

Open Discussion II: Future of the User Forum

- Given all these changes, how should we shape the future of the PRACE User Forum?

For instance:

- (How) can a User Forum work in a PRACE project that orients primarily towards user support?
- Would it be sensible to position the User Forum to support both PRACE and EuroHPC?
- If we choose to cater for both, what changes should we make in the User Forum Committee?
 - Dedicated liaisons for resp. PRACE and EuroHPC?
 - Any change of scope?

Thank you!

EuroHPC & PRACE

- EuroHPC will start delivering core-hours to researchers this year, and will supersede PRACE peak performance by a factor of 5 by 2022
- EuroHPC: >1.1 Exaflops peak performance, with 90-10 GPU-CPU split
- PRACE: 220 PFs peak performance, with 60-40 GPU-CPU split
- PRACE is still very relevant:
 - PRACE peer review process is the de facto gold standard
 - PRACE Training centres are generating the human capital that can use the European Data Infrastructure
 - PRACE High-level support teams are enabling new research groups to use Tier-0 resources efficiently
 - PRACE has entered new collaborations e.g. with the Fenix infrastructure to support analysis and data intensive science
 - Not all PRACE members are part of EuroHPC: UK is a notable

EuroHPC & PRACE

- EuroHPC will start delivering core-hours to researchers this year, and will supersede PRACE peak performance by a factor of 5 by 2022
- EuroHPC: >1.1 Exaflops peak performance, with 90-10 GPU-CPU split
- PRACE: 220 PFs peak performance, with 60-40 GPU-CPU split
- PRACE is still very relevant:
 - PRACE peer review process is the de facto gold standard
 - PRACE Training centres are generating the human capital that can use the European Data Infrastructure
 - PRACE High-level support teams are enabling new research groups to use Tier-0 resources efficiently
 - PRACE has entered new collaborations e.g. with the Fenix infrastructure to support analysis and data intensive science
 - Not all PRACE members are part of EuroHPC: UK is a notable

EuroHPC & PRACE

- EuroHPC will start delivering core-hours to researchers this year, and will supersede PRACE peak performance by a factor of 5 by 2022
- EuroHPC: >1.1 Exaflops peak performance, with 90-10 GPU-CPU split
- PRACE: 220 PFs peak performance, with 60-40 GPU-CPU split

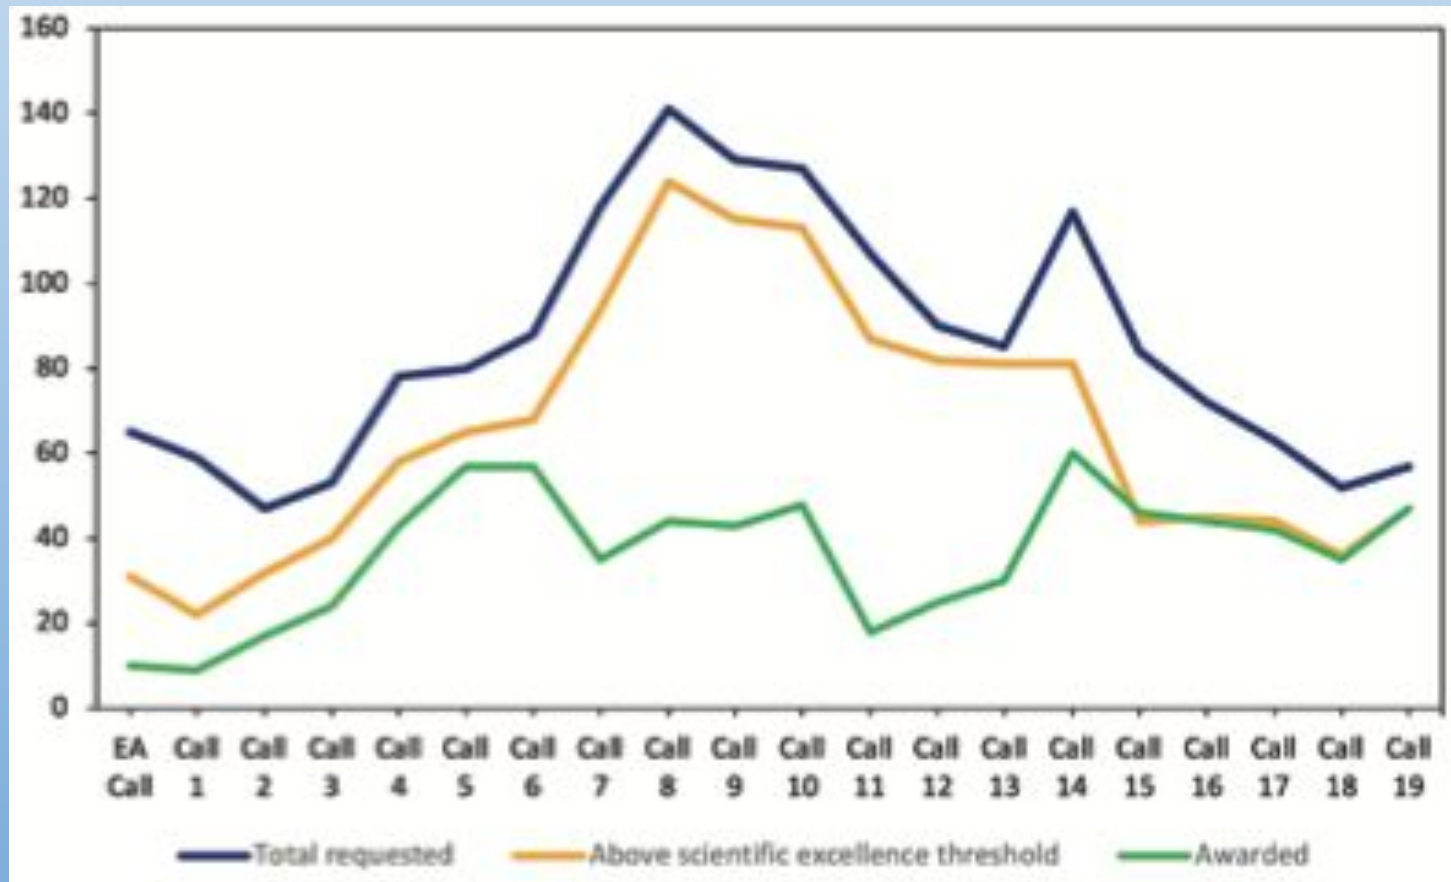
EuroHPC & PRACE

- EuroHPC will start delivering core-hours to researchers this year, and will supersede PRACE peak performance by a factor of 5 by 2022
- EuroHPC: >1.1 Exaflops peak performance, with 90-10 GPU-CPU split
- PRACE: 220 PFs peak performance, with 60-40 GPU-CPU split
- Our user survey indicates that majority of research groups cannot yet carry out their science using GPUs
- Some scientific applications run on GPUs, others will eventually, but an important subset does not match the underlying hardware
- PRACE has today a similar level of pure CPU performance compared to what EuroHPC will provide
 - PRACE delivers a wide variety of architectures and for a large fraction of users, the performance value will be similar to EuroHPC

EuroHPC & PRACE

- EuroHPC will start delivering core-hours to researchers this year, and will supersede PRACE peak performance by a factor of 5 by 2022
- EuroHPC: >1.1 Exaflops peak performance, with 90-10 GPU-CPU split
- PRACE: 220 PFs peak performance, with 60-40 GPU-CPU split
- The history of PRACE shows that
 - Maintaining a diversity of architectures is important to support a diversity of science applications
 - Input by the applied scientists from academia and industry, who are the users of the infrastructure, helps broadening the scope and improving the access rules.
- The User Forum believes that it would be constructive for EuroHPC to establish a similar user forum, to give users a way to channel feedback. This is common practice at all large science

Current usage of PRACE



EuroHPC & PRACE

- EuroHPC will start delivering core-hours to researchers this year, and will supersede PRACE peak performance by a factor of 5 by 2022
- EuroHPC: >1.1 Exaflops peak performance, with 90-10 GPU-CPU split
- PRACE: 220 PFs peak performance, with 60-40 GPU-CPU split
- Our user survey indicates that majority of research groups cannot yet carry out their science using GPUs
- Some scientific applications run on GPUs, others will eventually, but an important subset does not match the underlying hardware
- PRACE has today a similar level of pure CPU performance compared to what EuroHPC will provide

□ Broad usage of the EuroHPC capacity will require sustained high-level support or funding programs aimed at transitioning current Tier-0 research groups or communities to accelerated architectures. Otherwise there is a real risk of undersubscription

EuroHPC & PRACE

- EuroHPC will be available in 2022, this year, and will be scaled up to 1.5 exaFLOPS by 2022
- EuroHPC will be a CPU split
- PRACE will be a GPU split
- Our current hardware cannot
- Some hardware is available, but it is not as powerful as EuroHPC will be
- PRACE has today a similar level of pure CPU performance compared to what EuroHPC will provide

We need to maintain European excellence in software and application development!

Human resources are present throughout industry and academia, and software support is much less costly than hardware

□ Broad usage of the EuroHPC capacity will require sustained high-level support or funding programs aimed at transitioning current Tier-0 research groups or communities to accelerated architectures. Otherwise there is a real risk of undersubscription

Backup slides to be deleted

Activities and Contact

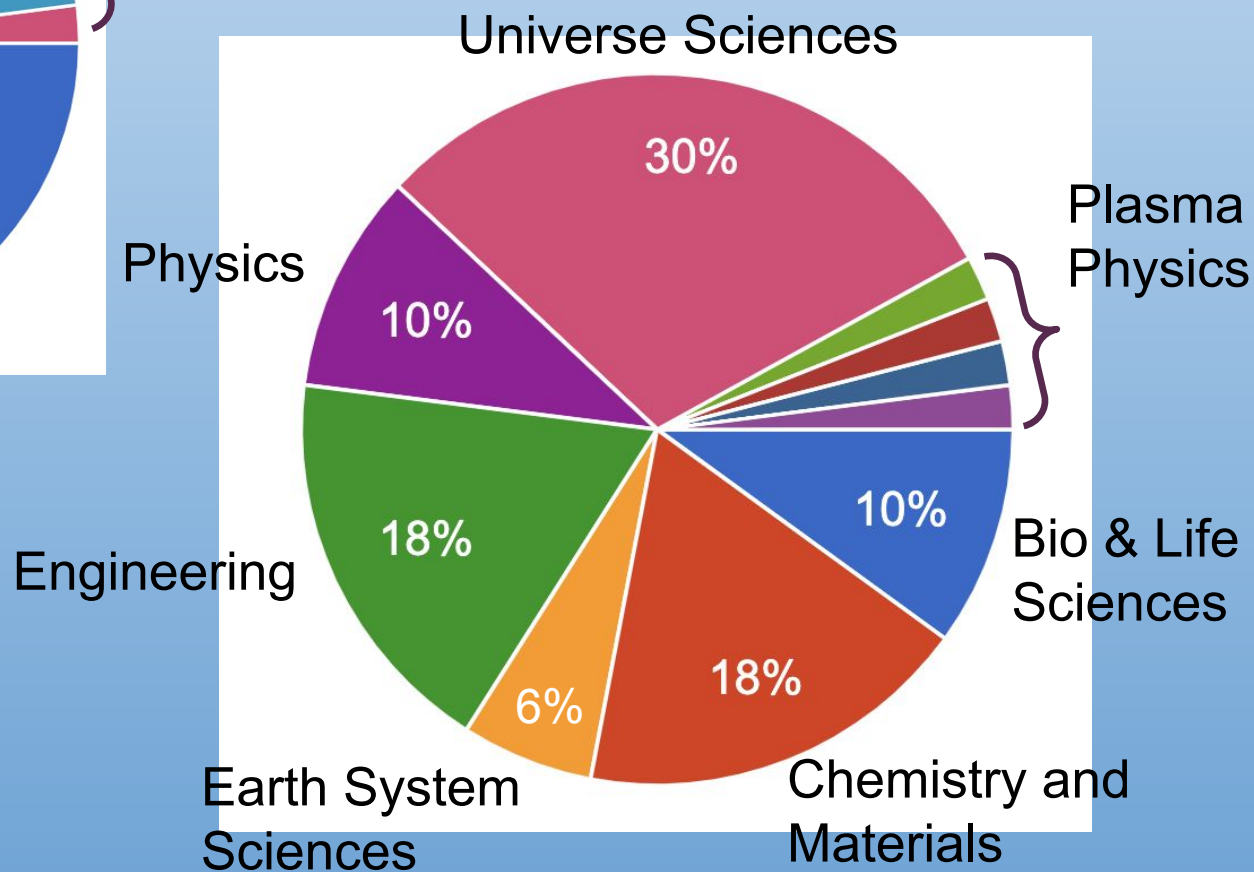

- General assembly at PRACE Days
- Outreach at conferences
 - Contents
 - ✓ contributed talks
 - ✓ PRACE representative
 - ✓ Discussion
- Contact: userforum@prace-userforum.eu

Surveying PRACE users in 2019

- Survey carried out in spring 2019 among the PIs of PRACE Tier-0 awards
- **Case:** understand HPC user requirements and their view of PRACE and HPC in Europe in general
- **Demographics:** gathered 50 responses from former and current Tier-0 PIs leading groups from a broad range of fields **representative**


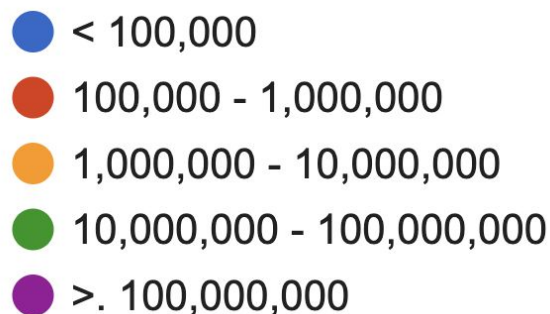
What is a typical Tier-0
HPC user?

What is a typical Tier-0 HPC user?




What is a typical Tier-0 HPC user?

Core hours used per year




- 10 – 100M CPUh / yr typical usage
- 12% use +100M / yr
- In house codes, or "standard" field specific codes

Type of code running on PRACE




What is a typical Tier-0 HPC user?

Architecture used for the codes



PRACE Machines Used



- Multi-core still most popular, but many-core getting uptake
- Even FPGAs!
- Users of all PRACE machines represented


What is a typical Tier-0 HPC user?

- Academic
- 10 – 100M CPUh / yr
- In house code
- Mostly multi-core optimised code, but increasingly also able to run on many-cores

PRACE Application Process

PRACE Application Process

Compared to other grants, how is the process of applying to PRACE?

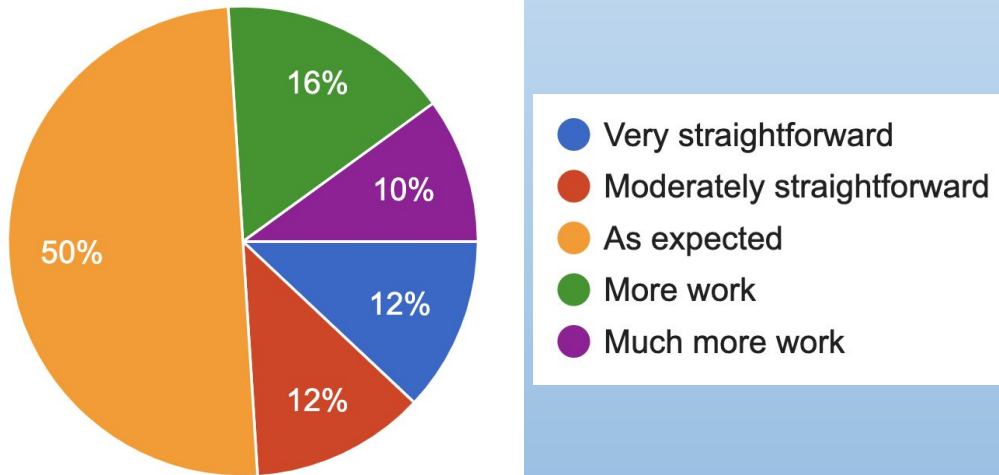


- Much less work
- Somewhat less work
- About the same
- More work
- Much more work

- Workload of PRACE Application process well balanced compared to other grants
- Important, because otherwise not enough applications


Applications: how easy to ...

Write Technical Justification



- Easier to write technical part than gathering the data

Gather data for Technical Justification




"More specific guidelines to answer the technical requirements"

"Different systems may require different scalability tests, not necessarily easy to perform during the application stage"


Applications: how easy to ...

Write Scientific Justification



- Very straightforward
- Moderately straightforward
- As expected
- More work
- Much more work

Respond to referee comments




- Very straightforward
- Moderately straightforward
- As expected
- More work
- Much more work

- Scientists know how to write their science!
- Referee system is working and not too much of a burden

PRACE Calls: frequency and size


Application Frequency and size



- 2 calls per year for 1 year allocations with smaller core hour minima
- 2 calls per year for 1 year large allocations (minimum 15 million core hours)
- 1 call per year for 1 year allocations with even larger minimum number of core hours (e.g. like INCITE)
- 2 year allocations
- 3 year allocations

- current design with larger Tier-0 allocations twice a year works.
- Users ok with re-location of grants to similar systems

Relevance of multiple systems




- Relevant to maximize the chances to be awarded and easy to include in our application
- Relevant to maximize the chances to be awarded but difficult to handle in our application
- Relevant because our project has different requirements matching different systems
- Not relevant or the added application burden is not worth the trouble

"Difficult timescale. Only large groups can live with this. For smaller groups turnaround too long compared to PhD / Postdoc contracts"

PRACE Infrastructure and Experience


PRACE Infrastructure & Experience

How well is the workflow supported




- Current machines accessible through PRACE supports the workflow of users very well
- Strongly geared towards large-scale capability modelling

Relevance of workflows



PRACE Infrastructure & Experience


How easy is it to start using the infrastructure?



- Very straightforward
- Moderately straightforward
- As expected
- Moderately difficult
- Very difficult

- Varying degrees of success with setting up access to HPC centres
- Reporting requirements (checks and balances!) reasonable


How does the reporting requirements compare to similar sized grants?



- Much less work
- Somewhat less work
- About the same
- More work
- Much more work

Bottom line – Infrastructure


Would you apply to PRACE again?



- Yes, definitely
- Yes, probably
- Maybe
- No, probably not
- No, definitely not

Bottom line – Infrastructure

Quotes from PRACE researchers



CAUTIONARY REMARKS

- “The obligation of using a minimum number of cores (1024) per task is not ideal”
- “The technical details in the new application form seems unreasonably detailed.”

PRAISE

- “I think in Europe there is no other alternative than PRACE if one needs allocation of size of order 100M core-hr”
- “Great opportunity to make progress in my research”
- “I really appreciate the 3 years project.”
- “The resources granted by PRACE are necessary for the completion of our work”

Is PRACE Superseded by other HPC facilities?

PRACE is the only option!

- No; PRACE is the only avenue via which we can apply for so many core hours.
- No. PRACE is the only option in Europe.
- No. We will need access to very large systems in the future that will not be available through other channels
- No, no, no, no, no,

PRACE alternatives

- National Tier-1 covers most. The application is easier (no need for technical details) but limited by shorter duration
- I feel that I am a valued customer of Jewels and SuperMuc. Get always feedback within very short time.
- NERSC. The machine is made for large scale computing.
- CSCS production

Is PRACE Superseded by other HPC facilities?

PRACE is the only option!

- No; PRACE is the only avenue via which we can apply for so many core hours
- No. PRACE
- No. We at
- will not
- No

PRACE is an indispensable infrastructure to perform modelling inaccessible on national resources

PRACE

- Na ed
- for
- I feel rMuc. Get
- always feel
- NERSC. The machine is in the scale computing.
- CSCS production

Sensible final comments by users

PRACE can impact the long-term HPC & science landscape

- “Suggestion for HLSTs: look for collaborations with HLST experts that result in new algorithms or implementations advancing the field.” [And maintain European excellence in applications!]
- “PRACE 3 year access is crucial for sustained long-term, ambitious projects from large groups.”

What have we learned?

- PRACE is a ***resounding success*** with the current users.
- System of application, review, reporting is fair and has a reasonable workload.
- ***Typical Tier-0 user***: 10 – 100 MCpuh / yr using in house developed code □ medium sized research groups with a large variety of codes are the typical Tier-0 users
- ***Application diversity*** is an important characteristic of European HPC landscape

contact: userforum@prace-userforum.eu

Future European HPC Landscape

- *Application diversity* is an important characteristic of European HPC landscape
- Preparing enough users to run on pre-Exascale systems requires diversity
- Many (most?) of the current “community codes” grew organically starting out as a code developed by a small team of researchers
- use the established mechanism of HLSTs to support development of (in-house) codes used currently on Tier-0 systems.