



PARTNERSHIP FOR
ADVANCED COMPUTING IN EUROPE



e-Infrastructure
Reflection Group



Federated Services and Data Management in PRACE

G. Erbacci, (CINECA and PRACE)

Collaborative Services for e-Infrastructure Commons
e-IRG Workshop, Bratislava 15-16 November 2016



PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Partnership for Advanced Computing in Europe

PRACE is an international not-for-profit association under Belgian law, with its seat in Brussels.

PRACE counts 25 members and 2 observers.

The **PRACE** Hosting Members are France, Germany, Italy and Spain.

PRACE is governed by the **PRACE** Council in which each member has a seat. The daily management of the association is delegated to the Board of Directors.

PRACE is funded by its members as well as through a series of implementation projects supported by the European Commission.





PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

4 Hosting Members offering core hours on 6 world-class machines + 1



JUQUEEN: IBM
BlueGene/Q
GAUSS/FZJ
Jülich, Germany



SuperMUC: IBM
GAUSS/LRZ
Garching, Germany



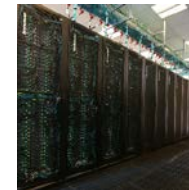
Hazel Hen: Cray
GAUSS/HLRS,
Stuttgart, Germany



CURIE: Bull Bullx
GENCI/CEA
Bruyères-le-Châtel,
France



MareNostrum: IBM
BSC, Barcelona,
Spain



Marconi: Lenovo
CINECA
Bologna, Italy



Piz Daint: Cray XC30
CSCS
Lugano, Switzerland

~30 PFlop/s in total

11.4 thousand million core hours awarded since 2010

+ Tier-1 Systems

26 PRACE national sites distributed in 19 different Countries are operational

> 17 PFlop/s in total

Apart **Tier-1 for Tier-0 services**, partners provide resources for **DECI calls**



PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE



465 scientific projects **enabled**



12.2 thousand million core hours awarded since 2010 with peer review, main criterion is **scientific excellence**. **Open R&D** access for **industrial users** with **>50 companies** supported



>7 350 people trained by **6 PRACE Advanced Training Centers** and others events



30 Pflop/s of peak performance on **7 world-class systems**



530 M€ of funding for **2010-2015**, access **free at the point of usage**



25 members, including **5 Hosting Members**
(France, Germany, Italy, Spain, Switzerland)



**Criterion:
Scientific
Excellence**

Access through PRACE Peer Review



Free-of-charge required to publish results at the end of the award period
Open to **international** projects



Project Access (*every 6 months*) **award period 1 to 3 years**
Individual researchers and groups
No restriction on nationality for both researcher and centre
Required to demonstrate technical feasibility of project



Preparatory Access (*cut-off date every 3 months*)
Optional support from PRACE experts
Prepare proposals for Project Access



PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Project Access



www.prace-ri.eu/call-announcements/



PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Preparatory Access Types

Type of Prep Access	Code Already Scales to Tier-0 Minimum ?	System	Duration	Apps enabling work needed for efficient running?	Expected outcome
Type A	Yes, this is a pre-requisite	Tier-0 only	3 months	No, scaling testing only	Application for Project Access for production runs
Type B	Yes, this is a pre-requisite	Tier-0 only	6 months	Yes, work will be done by researchers themselves	Application for Project Access for production runs
Type C	Yes, this is a pre-requisite	Tier-0 only	6 months	Yes, staffing resources from project requested	Application for Project Access for production runs
Type D	Possibly, but not necessarily. Potential must be demonstrated.	Tier-1 and Tier-0	12 months	Yes, staffing resources from project requested	Application to Prep Access A, B or C to prepare code for Tier-0 access, application for Project Access, SHAPE or DECI for production runs



PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE



- **SME HPC Adoption Programme in Europe**
- Programme to raise awareness and assist SMEs in taking advantage of HPC
 - Enable development of new products; reduce time-to-market and R&D costs; increase quality
- Aims to enable SMEs to overcome barriers to HPC adoption
 - Such as cost, lack of expertise, access to resources, ...
- Successful applicants get effort from a PRACE HPC expert and machine time at a PRACE centre

Seasonal Schools

- Have been running since 2008
- Offering top-quality face-to-face training events
- Organised around / all over Europe
- Topics range from generic intermediate to advanced
- From programming techniques to more specialised topical schools

Upcomming Seasonal Schools

- 27 - 30 September 2016
PRACE Autumn School 2016, Austria - Modern HPC Development for Scientists and Engineers
- 10 – 12 April 2017
PRACE Spring School 2017, Sweden - HPC in the Life Sciences





PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

PATC Programme 2016- 2017

- 79 courses, 215 training days
- New courses on forward-looking topics
 - New hardware and programming paradigms
 - Data science
- Collaboration with CoEs on several courses



PRACE Advanced Training Centres (PATCs)

BSC - Barcelona Supercomputing Center (Spain)
CSC - IT Center for Science (Finland)
CINECA - Consorzio Interuniversitario (Italy)

EPCC at the University of Edinburgh (UK)
GCS - Gauss Supercomputing Center (Germany)
MdS - Maison de la Simulation (France)

Programme: August 2016 to January 2017

August 2016

- Advanced OpenMP [EPCC]
- GPU Programming with CUDA [EPCC]

September 2016

- High Performance Molecular Dynamics [CINECA]
- Introduction to CINECA HPC System [CINECA]
- Introduction to High Performance Computing with C [CSC]
- Introduction to High Performance Computing with Fortran [CSC]
- Advanced MPI [EPCC]
- Object-Oriented Programming with Fortran [EPCC]
- Single node performance optimisation [EPCC]
- Advanced Fortran topics [GCS]

October 2016

- Parallel Programming Workshop [BSC]
- Python for computational science [CINECA]
- Data Intensive Analyses [CSC]
- Introduction to Parallel Programming [CSC]
- Practical Software Development [EPCC]
- Advanced Parallel Programming with MPI 3.1 [GCS]
- VI-HPS Tuning Workshop [GCS]
- Code optimization and debugging [MdS]

November 2016

- Earth science simulation environments [BSC]
- Debugging and Optimization of Scientific Applications [CINECA]
- HPC Methods for Computational Fluid Dynamics and Astrophysics [CINECA]
- Introduction to Parallel Computing with MPI and OpenMP [CINECA]
- Hybrid MPI/OpenMP Programming [MdS]

December 2016

- GPU Programming with OpenACC [CSC]
- Node-level Performance Engineering [GCS]
- Programming on GPUs [MdS]

January 2017

- Administration of Petaflop Machine [BSC]
- Material science codes on innovative HPC architectures: targeting exascale [CINECA]
- Python in High-Performance Computing [CSC]
- Efficient Parallel IO on ARCHER [EPCC]
- Introduction to hybrid programming in HPC [GCS]
- Advanced Visualization with Paraview [MdS]
- Parallel linear algebra [MdS]

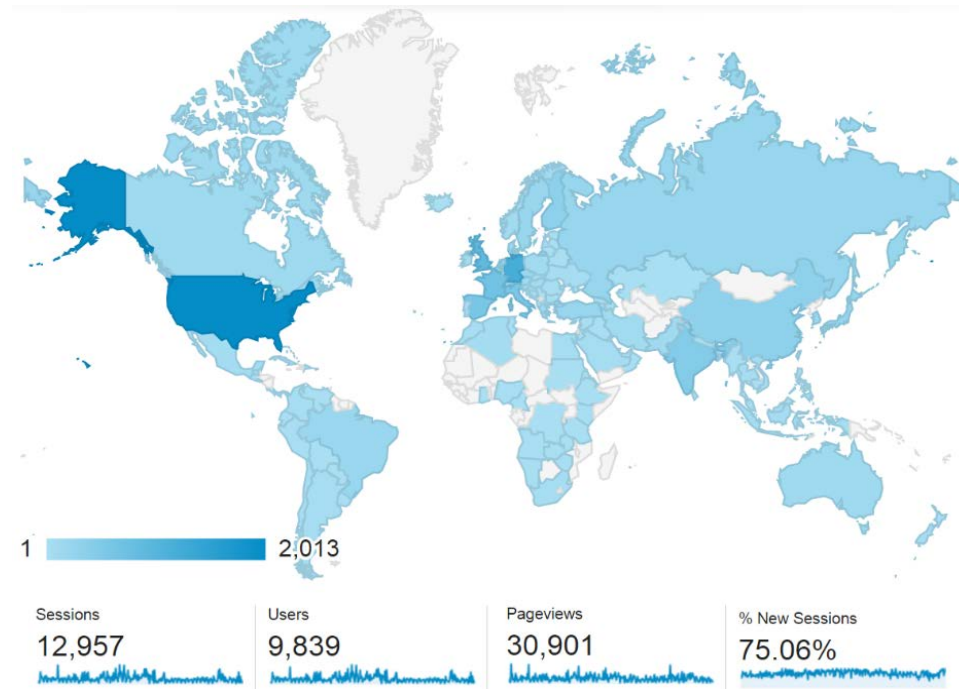
www.prace-ri.eu

The Implementation Phase of PRACE receives funding from the EU's Seventh Framework Programme (FP7/2007-2013) under grant agreement RI-312763 and from the EU's Horizon 2020 Research and Innovation Programme (2014-2020) under grant agreement 653838.



Training and Events Portal

- www.training.prace-ri.eu
- Single hub for the PRACE training events, training material and tutorials
- Number of page views increased by 25% and number of users by 45% since 2014



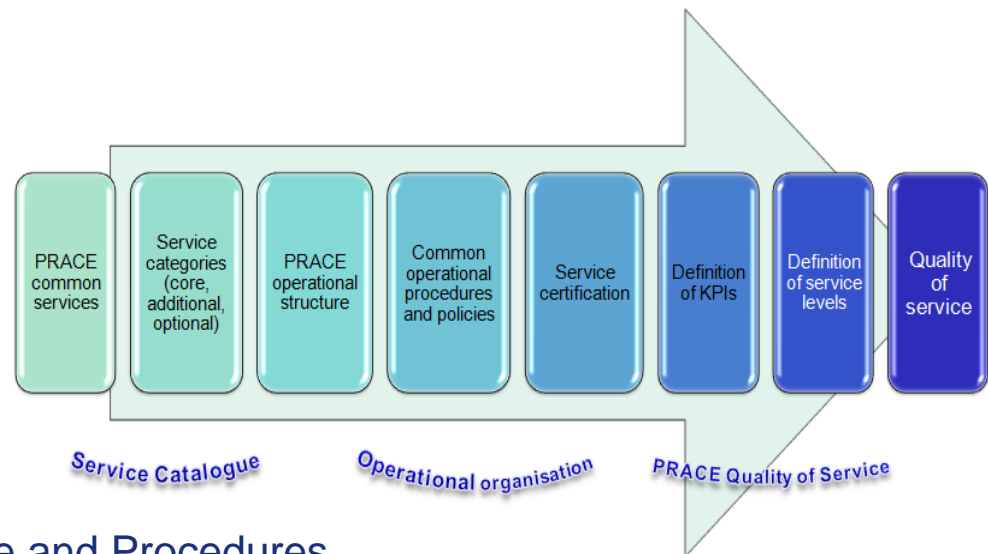


PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Operation and Coordination of the Comprehensive common PRACE Operational Services

- Common view of the PRACE infrastructure → more than a collection of individual systems
- Responsible for both Tier-0 systems and Tier-1 systems providing *Tier-1 for Tier-0* services

Key assets of the operational infrastructure



Infrastructure and common services

- Consolidated Operational Structure and Procedures
- PRACE Service Catalogue
- Implement operational Key Performance Indicators
- Security Forum to address security issues



PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

PRACE Operational Coordination Team

- Matrix organisation for Operations
 - Coordinated by WP6 Leader
 - Task Leaders for the deployment of service categories:
 - Networking, Data, Compute, AAA, User, Monitoring and Generic
- Site representatives are responsible for services at their site

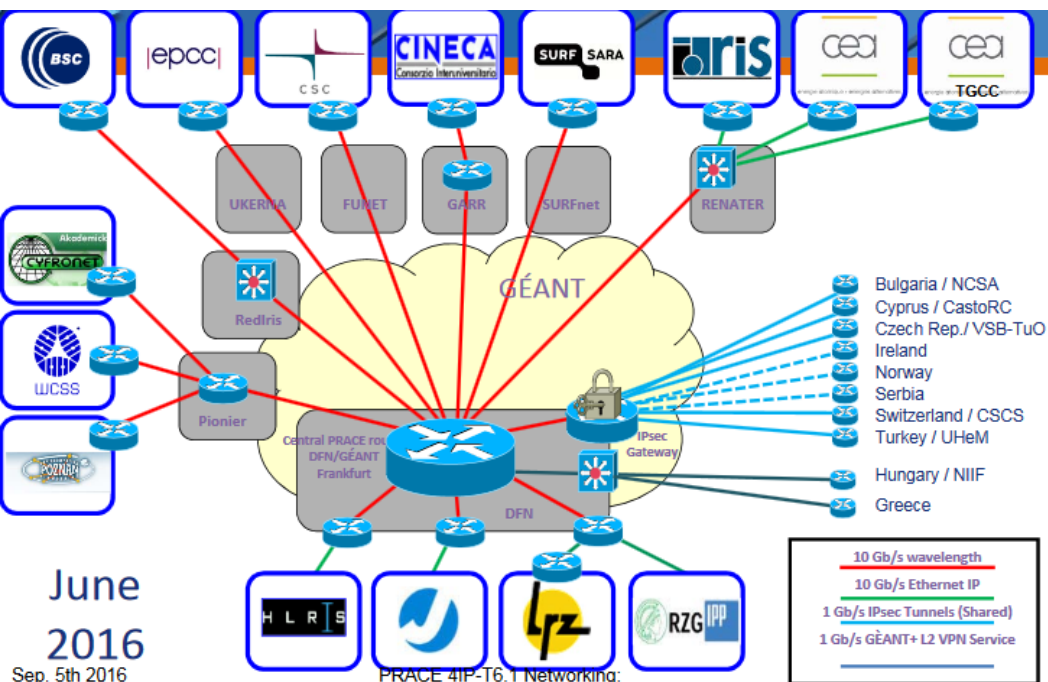
	Hosting partner 1 Site representative	Hosting partner 2 Site representative	Hosting partner N Site representative
Network Services Service category leader	Network services	Network services	Network services
Data Services Service category leader	Data services	Data services	Data services
Compute Services Service category leader	Compute services	Compute services	Compute services
AAA Services Service category leader	AAA services	AAA services	AAA services
User Services Service category leader	User services	User services	User services
Monitoring Services Service category leader	Monitoring services	Monitoring services	Monitoring services
Generic Services Service category leader	Generic services	Generic services	Generic services

- Bi-weekly telcos to discuss the status of services and sites and proposed or planned changes
- Changes are managed following a well defined procedure



PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Network services



Current PRACE dedicated network

A central L2/L3 switch in Frankfurt connecting

- 14 partners via 10 Gb/s wavelength

An IPSEC/GRE gateway in Frankfurt connecting

- 5 partners with 1 Gb/s IPSEC/GRE tunnels
- two partners via 1 Gb/s GÉANT-L2VPN connections

Future PRACE dedicated network

The infrastructure will be setup on the combined GÉANT / NRENs backbone providing a VPN between the PRACE partners (MDVPN service)

All partners will be connected by VLANs through their normal NREN connection to this PRACEVPN

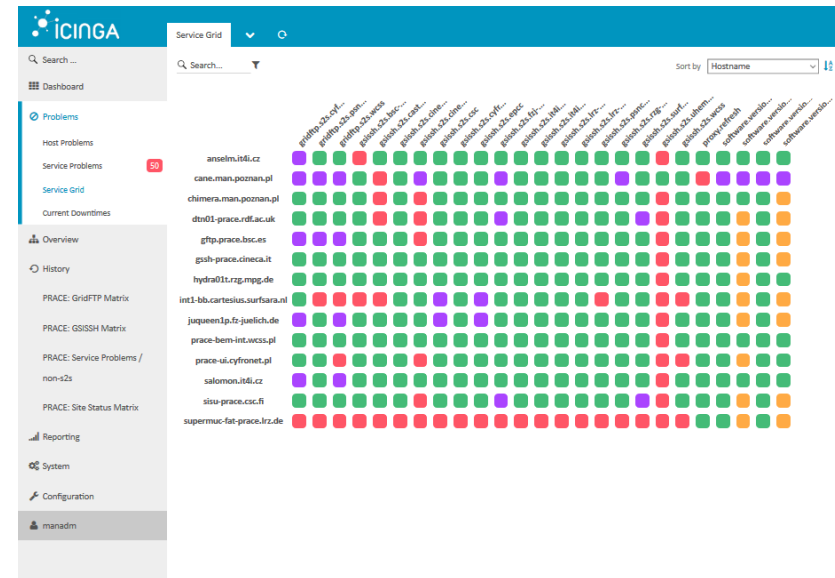
At NRENs, where MDVPN solutions are not available, partners can be connected via an MDVPN-Proxy provided by GÉANT



PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Monitoring of PRACE-RI Services

- INCA no more supported by SDSC
- Replaced by ICINGA 2 monitoring tools
- Deployed the new middleware and corresponding user interface for gathering and presenting monitoring data
- New domain name: <https://mon.prace-ri.eu/>
- 14 Hosts now connected
- 7 independent sets of services monitored
- Integration for all PRACE sites with valid user certificate



Checks:

software.version.libraries
software.version.compilers
software.version.shells

gssh.port – Host availability check based on gssh port state
gssh.s2s – site to site gssh connection check
gridftp.s2s – site to site gridftp connection check
software.version.tools





PRACE Security Forum

Coordinates security activities

- **Define Policies and Procedures:** to build “a trust model that allows smooth interoperation of the distributed PRACE services”
- **Risk reviews:** to define and maintain “An agreed list of software and protocols that are considered robust and secure enough to implement the minimal security requirements”
- **Operational security:** coordination of incident handling
- All PRACE operational partners are members of the Security Forum
- Collaboration with other large distributed computing infrastructures (EGI, EUDAT, XSEDE, WLCG, OSG) on policies and procedures
- Continues the representation of PRACE as relying party of EUGridPMA, the policy authority for trusted Certificate Authorities



Security collaboration

- Operational Security
 - Collaboration with EGI CSIRT and EUDAT on sharing of information on incidents and vulnerabilities
 - Accreditation of PRACE CSIRT team at Trusted Introducer service from GEANT ongoing
- AAI
 - PRACE is Relying Party of EUGridPMA, the policy authority for the distribution of trusted Certificate Authorities
 - PRACE is represented in the AARC (Authentication and Authorisation for Research and Collaboration) Project <https://aarc-project.eu>
 - AARC objective: Enable the use of existing user credentials by the federation of existing Identity Providers and Service Providers
- **WISE Information Security for Collaborating E-Infrastructures**
 - A trusted global framework where security experts can share information on different topics like risk management, experiences about certification process and threat intelligence
 - Joint effort of GEANT SIG-ISM and SCI (EGI, EUDAT, HBP, PRACE, WLCG, XSEDE)



Analysis and Development of Prototypal New Services

- Provision of urgent computing services
- Link with large-scale scientific instruments
 - Link with the European Synchrotron Radiation Facility (CaSToRC)
 - MIC-oriented Multithreading for HEP and Health Geant4 Computations (NCSA)
 - HPC support for Extreme Light Infrastructure ELI-ALPS project (NIIF)
 - Linking Next Generation Sequencers with PRACE (UiO)
 - Large Synoptic Survey Telescope (CNRS)
- Smart post processing tools including in-situ visualisation
- Provision of repositories for European open source scientific libraries and applications, to promote wide adoption of European products
 - Analyse and investigate the prototypal implementations at the pre-production level (involving first Tier-1 systems and then Tier-0 systems)
 - Investigate the possible adoption in a next phase as production services



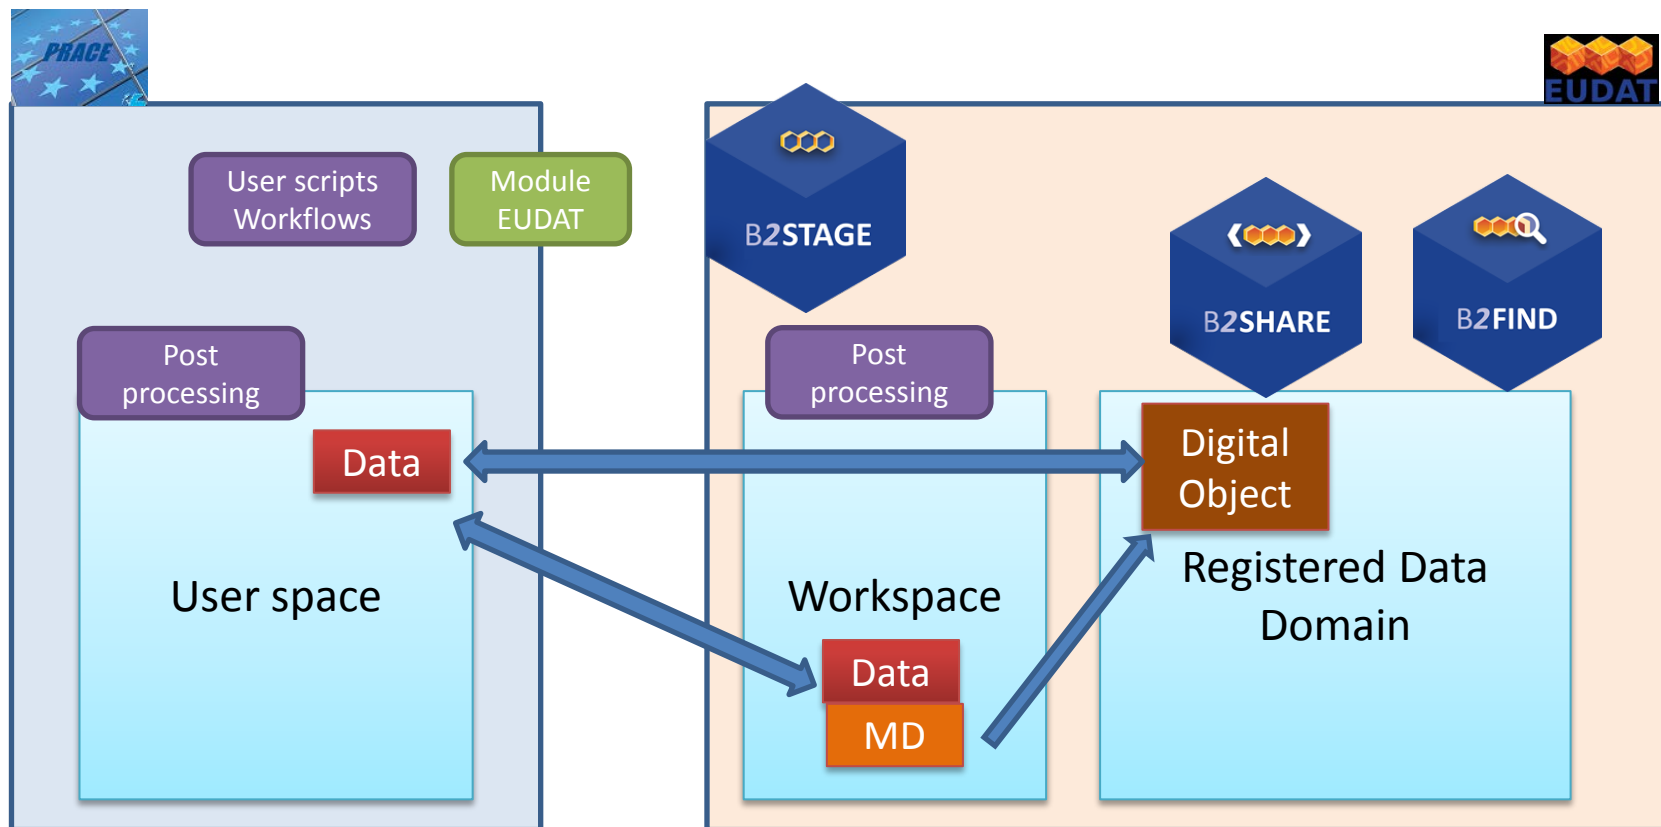
Data Collaboration with EUDAT

- MoU signed between PRACE and EUDAT
- Data pilots proposal analysis to identify use cases
- Get required changes on EUDAT services roadmap
- Data pilot identification:
 - **DECI Call 13: 5 pilots**
 - Deliver data management training to data pilots team
 - ongoing work to make it available for PRACE users
 - Gather detailed requirements
 - Resources available for the project
 - Detailed timelines
 - Data Management Plan
 - Technical constraints
 - Implementation ongoing in close collaboration with EUDAT team



- Contacts with CoEs on Data Management Issues

Overview of the 4 pilots





PRACE – GEANT Collaboration

- MoU signed between PRACE and GEANT
- Legal fundament to collaboration between both projects
- The current PRACE network infrastructure is spanning across several NRENs and the GÉANT infrastructure since more than 10 years
- A continuous adjustment to PRACE requirements but also to GÉANT and NREN services is essential to **improve the performance and transfer capacity** as well as to **secure the communication** streams across a global European network
- GEANT is interested in testing its **upcoming services**, if those fit user requirements, the same applies to PRACE where **evaluating new services** in an early state might be very valuable.



PRACE-EGI Collaboration

- MoU signed between PRACE and EGI Engage Projects
- Better identifying and coordinating common actions **toward the interoperability** and the **cross-adoption of different resources**
- Collaborating to define **common services** and defining an integrated catalogue of e-Infrastructure services, adopted from research communities
- Defining **common actions on security**, participating to the common boards defined by the WISE series of conferences;
- Providing a **joint training actions** on security aspects





Federated Services and Data Management

- Toward the interoperability with other public and private e-Infrastructures, lowering technical and policy barriers by piloting concrete use cases with user communities
- Is important to provide European researchers and industries a seamless access to HPC and data resources and services
 - Combination of respective services catalogue
 - Harmonization of access policy fostering the uptake of services on the long-term
 - Security for Collaboration among Infrastructures
- Implement the Open Science vision where resources of any kind and size are accessible without any technical barrier