EuroHPC

Andrej Filipčič

HPC Plans

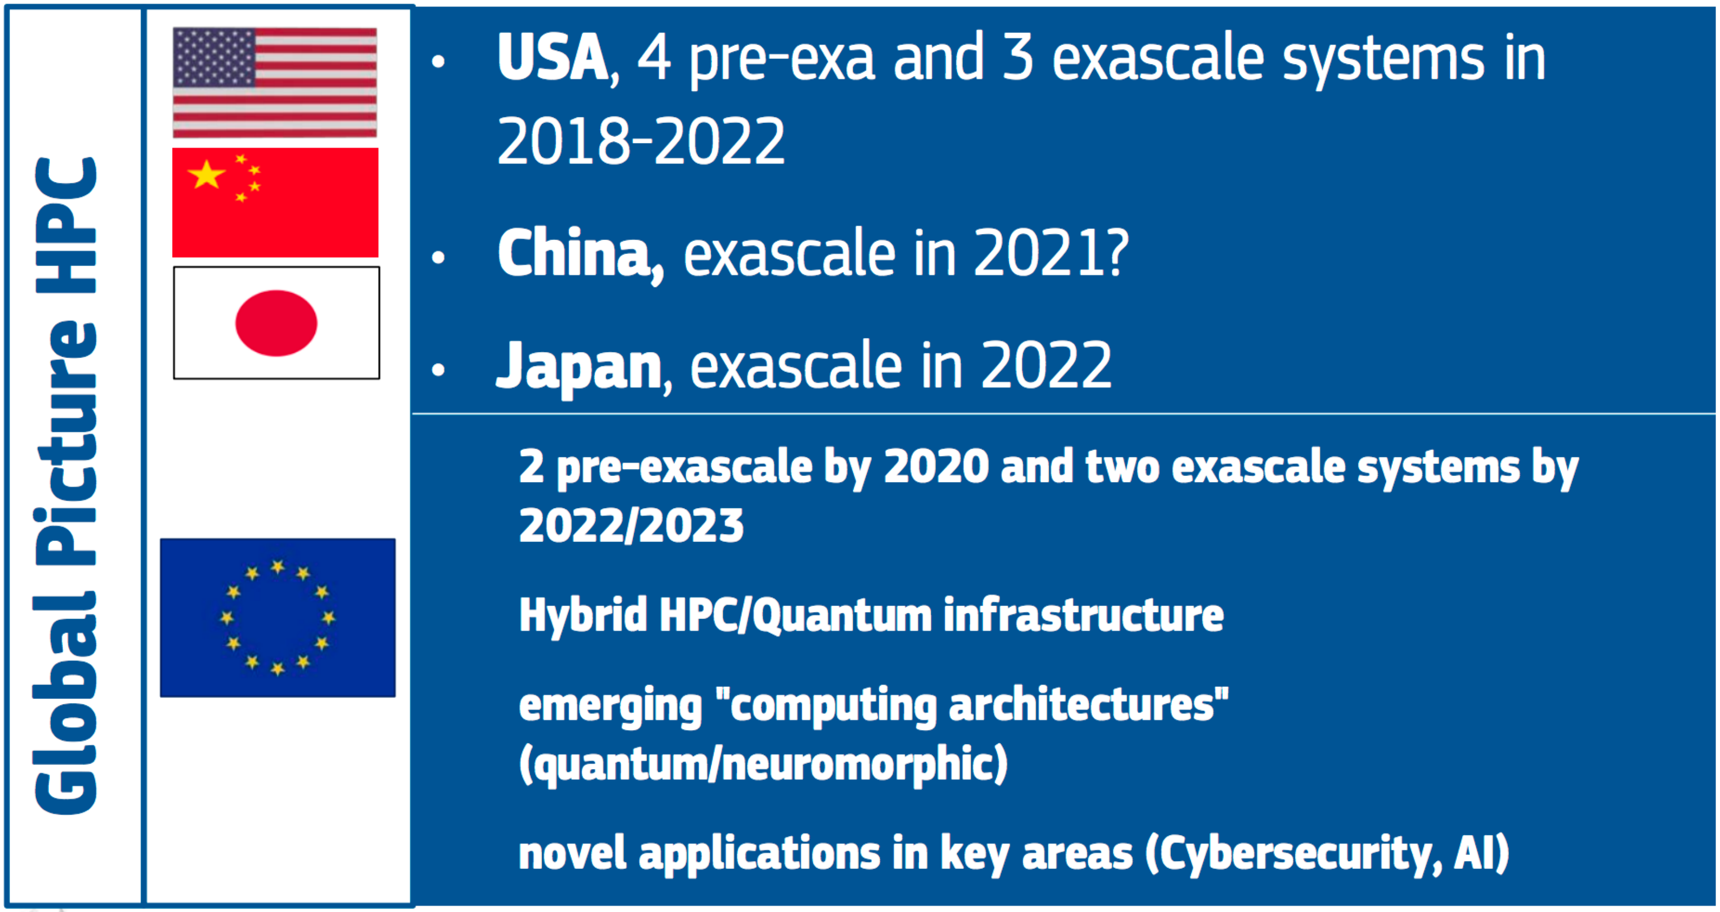
Global Picture HPC
• USA, 4 pre-exa and 3 exascale systems in 2018-2022
• China, exascale in 2021?
• Japan, exascale in 2022
2 pre-exascale by 2020 and two exascale systems by 2022/2023
Hybrid HPC/Quantum infrastructure
emerging "computing architectures" (quantum/neuromorphic)
novel applications in key areas (Cybersecurity, AI)

20 Member States, atm

#EuroHPC (High Performance Computing)
Declaration
Signatory European countries

Seven countries – France, Germany, Italy, Luxembourg, Netherlands, Portugal and Spain – signed the declaration in March 2017.

Since then, another thirteen countries – Belgium, Slovenia, Bulgaria, Switzerland, Greece, Croatia, Czech Republic, Cyprus, Poland, Lithuania, Austria, Finland, and Sweden – have also signed.

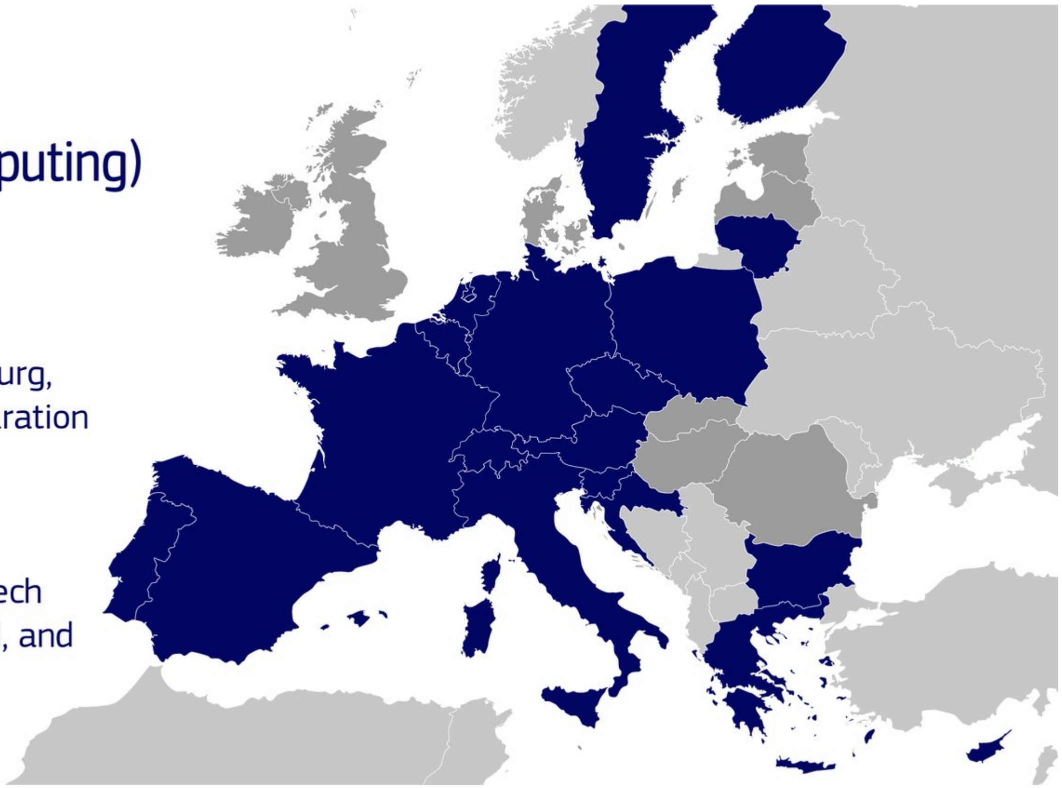

Timeline

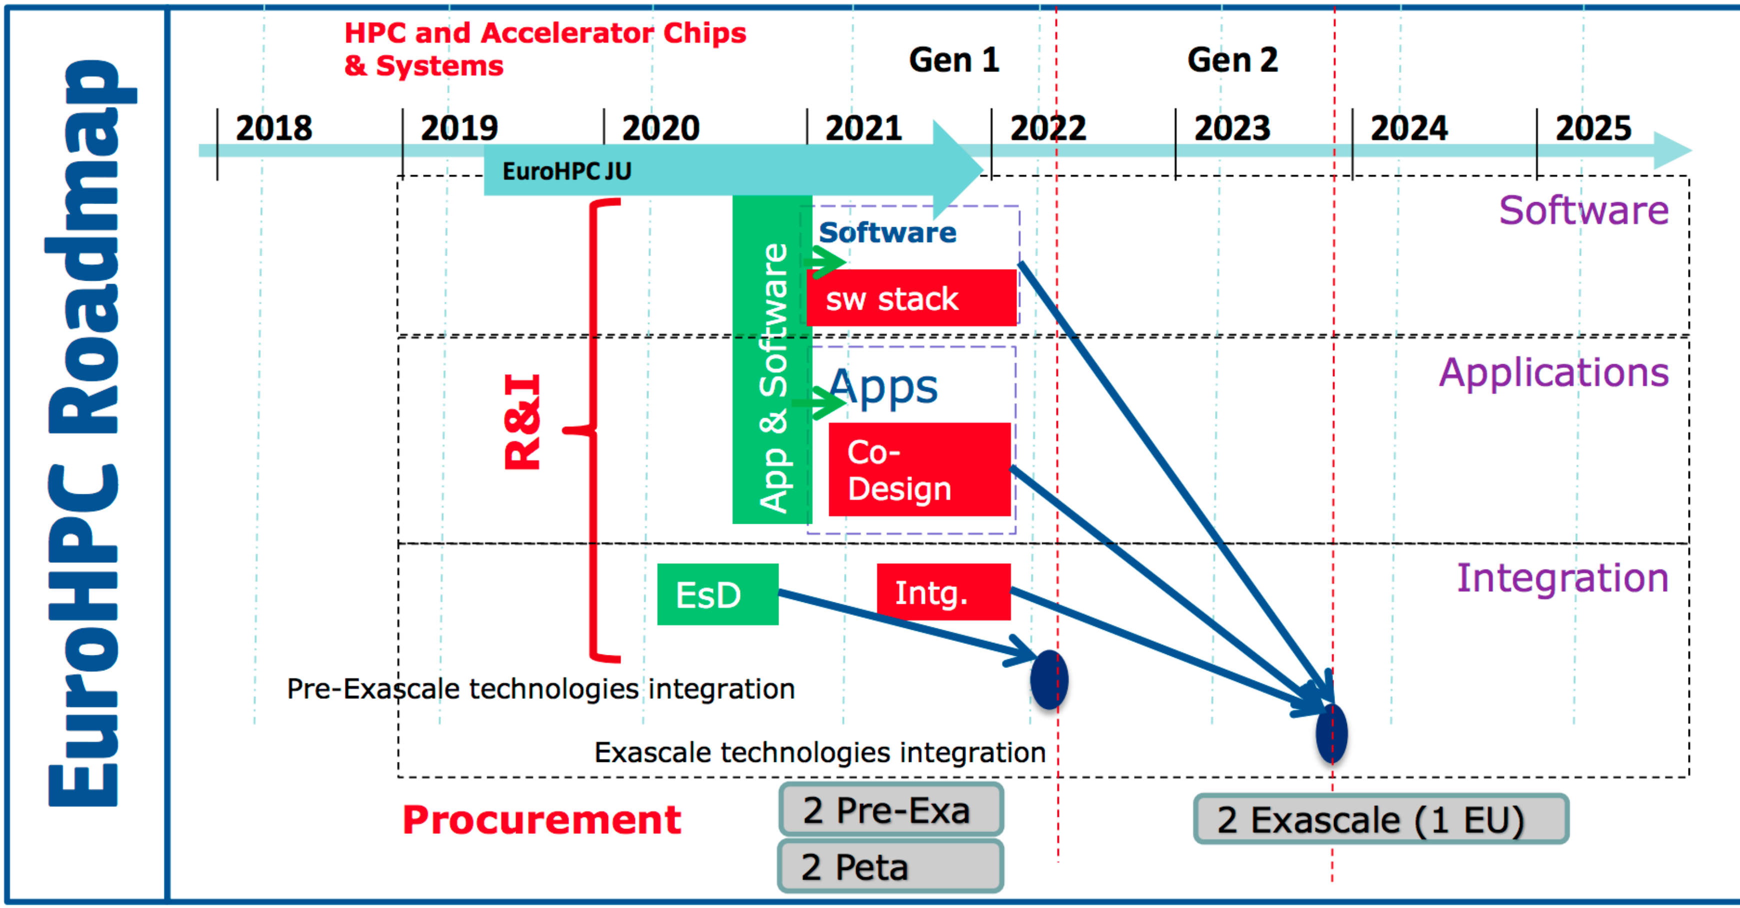

EuroHPC Roadmap
HPC and Accelerator Chips
& Systems
Gen 1
Gen 2
2018
2019
2020
2021
2022
2023
2024
2025
EuroHPC JU
Software
Software
sw stack
App & Software
R&I
Apps
Co-
Design
Applications
EsD
Intg.
Integration
Pre-Exascale technologies integration
Exascale technologies integration
Procurement
2 Pre-Exa
2 Peta
2 Exascale (1 EU)

Funds

- Effort started early 2017 at the EU level
- Plan: 2018-2020
  - Funding of ~1.5B euro (0.5B from EC) in 2018-2020
  - Build 2 pre-exascale machines by 2020(21) - ~100-200 PFlops
  - Build 2-4 peta-scale machines
- Plan: 2021-2026
  - Build 2 exascale machines by 2024(25) - ~1EFlops
  - Proposed Funding 2.7B euro (out of 9.2B for digital infrastructure)

BDV BIG DATA VALUE

HPC, Big Data and Deep Learning

		Supercomputing (SC)	Deep Learning (DL)	Big Data (BD)
Apps...	Boundary Interaction Services	**In-house, commercial & OSS applications** [e.g. Paraview], **Remote desktop** [e.g. Virtual Network Computing (VNC)], **Secure Sockets Layer** [e.g. SSL certificates]	**Framework-dependent applications** [e.g. NLP, voice, image], **Web mechanisms** [e.g. Google & Amazon Web Services], **Secure Sockets Layer** [e.g. SSL certificates]	**Framework-dependent applications** [e.g. 2/3/4-D], **Secure Sockets Layer** [e.g. SSL certificates]
Middleware & MGMT	Processing Services	**Domain specific frameworks** [e.g. PETSc], **Batch processing of large tightly coordinated parallel jobs** [100s – 10000s of processes communicating frequently with each other]	**DNN training & inference frameworks** [e.g. Caffe, Tensorflow, Theano, Torch], **DNN numerical libraries** [e.g. dense LA], **DNNs, statistics, diagnostics** [e.g. ?]	**Machine Learning** (traditional) [e.g. Mahout, Scikit-learn], **Analytics / Statistics** [e.g. Python, ROOT, R, Matlab, SAS, SPSS, Sci-Py], **Iterative** [e.g. Apache Hama], **Interactive** [e.g. Dremel, Drill, Tez, Impala, Shark, Presto, BlinkDB], **Batch / Map-Reduce** [e.g. MapReduce, YARN, Sqoop, Spark], **Real-time/streaming** [e.g. Flint, YARN, Druid, Pinot, Storm, Samza, Spark streaming]
	Model / Information Management Services	**Data Storage**: Parallel File Systems [e.g. Lustre, GPFS, BeeGFS, PanFS, PVFS], **I/O libraries** [e.g. HDF5, PnetCDF, ADIOS]	**Data Storage** [e.g. HDFS, Hbase, Amazon S3, GlusterFS, Cassandra, MongoBD, Hana, Vora]	**Serialization** [e.g. Avro], **Meta Data** [e.g. HCatalog], **Data Ingestion & Integration** [e.g. Flume, Sqoop, Apache Nifi, Elastic Logstash, Kafka, Talend, Pentaho], **Data Storage** [e.g. HDFS, Hbase, Amazon S3, GlusterFS, Cassandra, MongoBD, Hana, Vora], **Cluster Mgmt** [e.g. YARN, MESOS]
	Communication Services	**Messaging & Coordination** [e.g. MPI/PGAS, direct fabric access], **Threading** [e.g. OpenMP, task-based models]	**Messaging & Coordination** [e.g. Machine Learning Scaling Library (MLSL)]	**Messaging** [e.g. Apache Kafka (streaming)]
	Workflow / Task Services	**Conventional compiled languages** [e.g. C/C++/Fortran], **Scripting languages** [e.g. Python, Julia,]	**Scripting languages** [e.g. Python]	**Workflow & Scheduling** [e.g. Oozie], **Scripting languages** [e.g. Keras, Mocha, Pig, Hive, JAQL, Python, Java, Scala]
System SW	System Management & Security Services	**Domain numerical libraries** [e.g. PETSc, ScaLAPACK, BLAS, FFTW, ...], **Performance & debugging** [e.g. DDT, Vtune, Vampir], **Accelerator APIs** [e.g. CUDA, OpenCL, OpenACC] **Data Protection** [e.g. System AAA, OS/PFS file access control] **Batch scheduling** [e.g. SLURM], **Cluster management** [e.g. OpenHPC], **Container Virtualization** [e.g. Docker], **Operating System** [e.g. Linux OS Variant]	**Batching for training** [built into DL frameworks], **Reduced precision** [e.g. Inference engines], **Load distribution layer** [e.g. Round robin/load balancing for inference], **Accelerator APIs** [e.g. CUDA, OpenCL], **LA numerical libraries** [e.g. BLAS, LAPACK, ...] **Virtualisation** [e.g. Dockers, Kubernetes, VMware, Xen, KVM, HyperX], **Operating System** [e.g. Linux (RedHat, Ubuntu, etc.), (Windows?)]	**Distributed Coordination** [e.g. ZooKeeper, Chubby, Paxos], **Provisioning, Managing & Monitoring** [e.g. Ambari, Whirr, BigTop, Chukwa], **SVM systems** [e.g. Google Sofia, libSVM, svm-py, ...], **Hardware Optimization Libraries** [e.g. cuDNN, MKL-DNN, etc.)] **Virtualisation** [e.g. Dockers, Kubernetes, VMware, Xen, KVM, HyperX], **Operating System** [e.g. Linux (RedHat, Ubuntu, etc.), Windows]
Hardware	Infrastructure	**Local storage** [e.g. Storage & I/O nodes, NAS]; **Servers** [e.g. CPU & Memory) [Gen Purpose CPU nodes, GPUs, FPGAs]; **Network** [e.g. Infiniband & OPA fabrics]	**Local storage** [e.g. Local storage or NAS/SAN]; **Servers** [e.g. CPU & Memory) [Gen Purpose CPU + GPU/FPGA, TPU]; **Network** [e.g. Ethernet‡]	**Local storage** [e.g. Direct attached Storage]; **Servers** [e.g. CPU & Memory, [Gen Purpose CPU hyper- convergent nodes]; **Network** [e.g. Ethernet fabrics]

Source: "Creating synergies across HPC & Big Data platforms", BDVA-ETP4HPC White Paper

ETP4HPC - Strategic Research Agenda

2.1
THE VALUE OF HPC

2.1.1
HPC as a Scientific Tool

Scientists from throughout Europe increasingly rely on HPC resources to carry out advanced research in nearly all disciplines. European scientists play a vital role in HPC-enabled scientific endeavours of global importance, including, for example, CERN (European Organisation for Nuclear Research), IPCC (Intergovernmental Panel on Climate Change), ITER (fusion energy research collaboration), and the newer Square Kilometre Array (SKA) initiative. The PRACE Scientific Case for HPC in Europe 2012 – 2020 [PRACE] lists the important scientific fields where progress is impossible without the use of HPC.

http://www.etp4hpc.eu/sra-2017.html

- CERN is one of top EU scientific endeavours
- But in the document, HEP requirements are not mentioned, apart from lattice QCD
- Future EU HPC centers will be extended to data processing facilities – eg ESiWACE needs large storage, transfers and remote processing (distributed systems)
- Most intensive-computing communities are participating in EuroHPC, HEP is left out for now
- Increased funding from both EC and Member state can result in lower funding of dedicated WLCG infrastructure
- It should be discussed how to ensure HEP presence in future design of EuroHPC landscape

How to proceed

- WLCG strategy document as the underlying source of information
- WLCG could/should nominate a representative on EuroHPC User requirements working group to represent the WLCG and wider HEP interests in EuroHPC (maybe even PRACE)
- Presentation on the next WG meeting with key points to address the WLCG/CERN/HEP needs
  - Many of HEP specific requirements have been explicitly stated by other communities/CoEs, eg HTC/HPC convergence, data intensive workflows, remote processing and analytics, data transfers and high-throughput networks, permanent storage, archival and data preservation
  - Additional material/presentations can be provided
- Balance between HPC/EuroHPC and WLCG infrastructure needs to be carefully considered. Even future HPCs cannot address all the needs of WLCG (eg definitely not the data lakes)
  - Eg: in Slovenia, a 100k-core national HPC will be built by 2020, Tier-2 will mostly invest in storage and less in CPU.

Backup

Organizational Structure

- Joint Undertaking (JU)
- Member States (Participating States)
- Hosting Entity: country or consortium of states hosting the center
- Working groups
  - Hosting and Procurement
  - User requirements
- Funding resources for pre-exascale
  - ~500M EU funding
  - ~500M MS funding
  - ~400M “private” funding, operational costs…

Relation to PRACE and other communities

- PRACE plans PRACE3 after 2020
- Relation to EuroHPC not quite clear yet but:
  - Likely adoption of PRACE model for access
  - Level of cooperation and integration with JU under discussion
- Centers of Excellence
  - PPI4HPC, POP, CompBioMed, ESiWACE,...

EDI & PRACE

- Offer a variety of system architectures to enable traditional scientific simulations
  - General purpose HPC centers and topical centers
- Addressing the convergence of HPC, HTC, HPDA and AI
  - FENIX and EUDAT projects - data management
- How to handle the large volume of data generated
  - Generated centrally at large RIs or distributed (e.g. sensor networks)
  - Rethink data movement (from edge to the data center) :
    - Local (at data source) processing facilities for data reduction
    - Central or distributed storage
    - How to support end to end workflows
- How to provide HPC/HTC capacity to large scale scientific instruments
- Need for an even tighter coordination of national infrastructure procurements in EU

Extend current PRACE activities to EDI

- PRACE as main contributor to EDI with GÉANT
- Access to JU- & Tier-0-infra provided by HMs for Open R&D for science & industry
- Offer training, code enabling, communication, Tier-1 for Tier-0 services provided by PRACE partners & AISBL office
- Extend services towards industry (SHAPE-Fortissimo), and to the public sector
- Local support to Tier-2 services across Europe

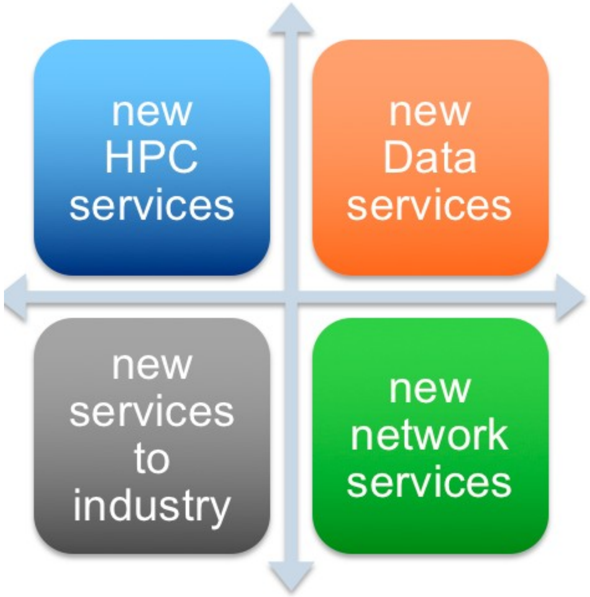

- ➢ New extended services provided by partners
- ➢ EDI as a one-stop-shop for all EU project/infra on HPC and data
- ➢ EOSC as potential vehicle to offer services to industry and public

PRACE, EDI, EuroHPC landscape

In PRACE (1,2)

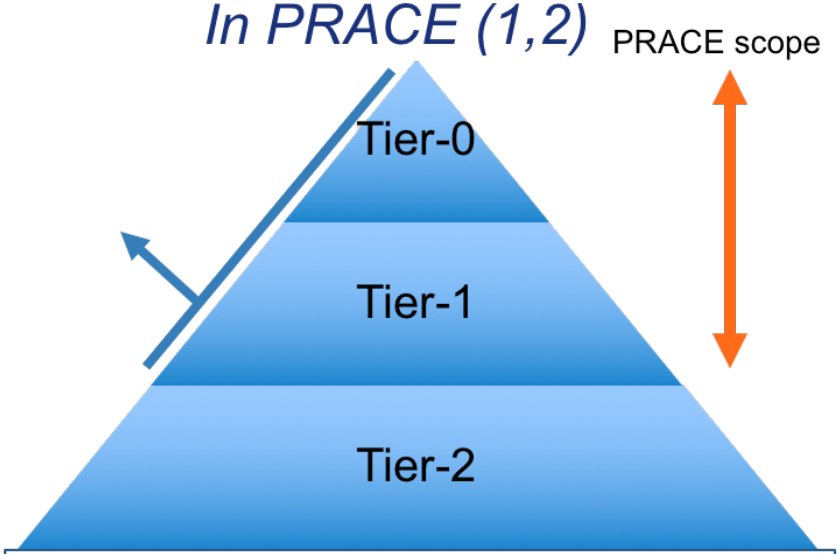

PRACE providing core services:

- Peer Review access to systems
- Training
- Code enabling (HLST)
- Communication / dissemination
- New communites engagement (SHAPE, …)
- With support from MS and EC

In EDI with EuroHPC

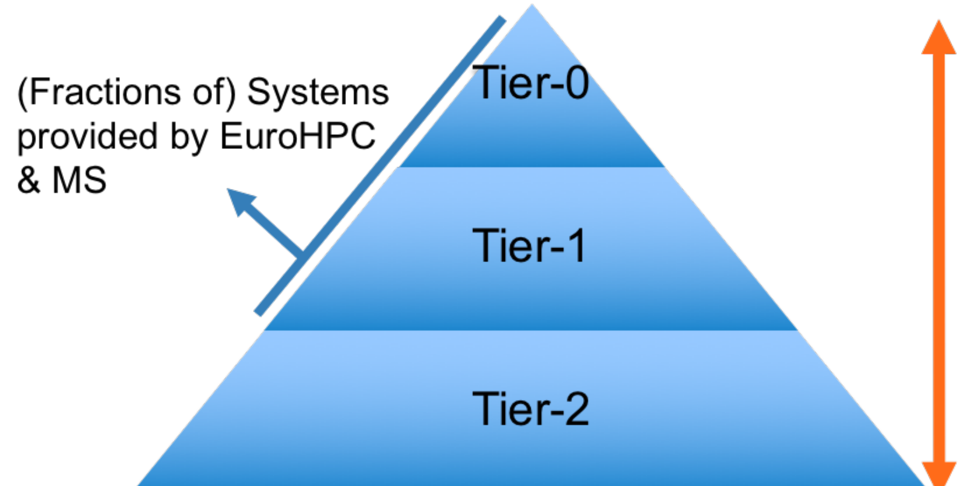

PRACE to provide same core services and propose additional (paid) services:

- New services for industry
- Tight link with scientific instruments / EC projects
- Urgent computing, …
- With support from MS and EuroHPC