IT4Innovations
National supercomputing center

Overview

Vít Vondrák




cesnet

Prague, January 29th-30th, 2019

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

IT4Innovations infrastructure

1. Anselm
Rpeak 94TFlop/s
Rmax 73TFlop/s
from June 2013

- 209 compute nodes
- 3344 Intel Sandy bridge cores
- 15136 GB RAM (64, 96, 512)
- 24 nVidia Tesla K20
- 4 Intel Xeon Phi (KNC)

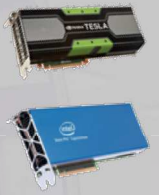

2. IT4I building
500m2 computer room
2x 2,5MVA power supply
July 2014

3. Salomon
Rpeak 2 Pflop/s
Rmax 1,5 Pflop/s
from July 2015

- 1008 compute nodes
- 24192 Intel Haswell cores
- 129024 GB RAM (128TB)
- 864 Intel Xeon Phi (KNC)

Kilo (10^3) Mega (10^6) Giga (10^9) Tera (10^{12}) Peta (10^{15}) Exa (10^{18})

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

Salomon in Top500.org

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
35	EPSRC/University of Edinburgh United Kingdom	**ARCHER** - Cray XC30, Intel Xeon E5 v2 12C 2.700GHz, Aries interconnect Cray Inc.	118,080	1,642.5	2,550.5	3,306
36	Grand Equipement National de Calcul Intensif - Centre Informatique National de l'Enseignement Suprieur (GENCI-CINES) France	**Occigen** - bullx DLC, Xeon E5-2690v3 12C 2.6GHz, Infiniband FDR Bull, Atos Group	50,544	1,628.8	2,102.6	935
37	IBM Development Engineering United States	Power 775, POWER7 8C 3.836GHz, Custom Interconnect IBM	62,944	1,587.0	1,931.6	3,576
38	ECMWF United Kingdom	Cray XC30, Intel Xeon E5-2697v2 12C 2.7GHz, Aries interconnect Cray Inc.	83,160	1,552.0	1,796.3	
39	ECMWF United Kingdom	Cray XC30, Intel Xeon E5-2697v2 12C 2.7GHz, Aries interconnect Cray Inc.	83,160	1,552.0	1,796.3	
40	IT4Innovations National Supercomputing Center, VSB-Technical University of Ostrava Czech Republic	**Salomon** - SGI ICE X, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR, Intel Xeon Phi 7120P HPE	76,896	1,457.7	2,011.6	1,538
41	Science and Technology Facilities Council - Daresbury Laboratory United Kingdom	**Blue Joule** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	131,072	1,431.1	1,677.7	657
42	Air Force Research Laboratory United States	**Spirit** - SGI ICE X, Xeon E5-2670 8C 2.600GHz, Infiniband FDR HPE	73,584	1,415.5	1,530.5	1,606
43	KTH - Royal Institute of Technology Sweden	**Beskow** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	53,632	1,397.0	1,973.7	786
44	CEA/TGCC-GENCI France	**Curie thin nodes** - Bullx B510, Xeon E5-2680 8C 2.700GHz, Infiniband QDR Bull, Atos Group	77,184	1,359.0	1,667.2	2,132

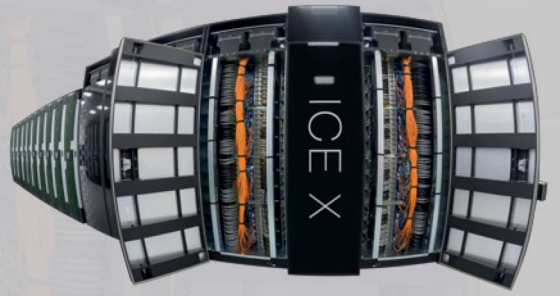

List	Highest Rank	Systems
2018/11	214	1
2018/06	139	1
2017/11	87	1
2017/06	78	1
2016/11	67	1
2016/06	55	1
2015/11	47	1
2015/06	39	1

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

Infrastructure: plan 2019-2020

EUROPEAN UNION
European Structural and Investment Funds
Operational Programme Research,
Development and Education

MINISTRY OF EDUCATION, YOUTH AND SPORTS

IT4Innovations National Supercomputing Centre – Path to Exascale

Anselm

May 2019

- Small cluster II
- 189 nodes, 840TF
- GPU acceleration
- € 3 mil.

Salomon

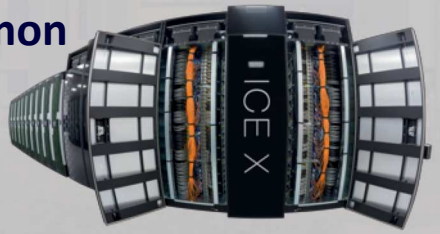

2020

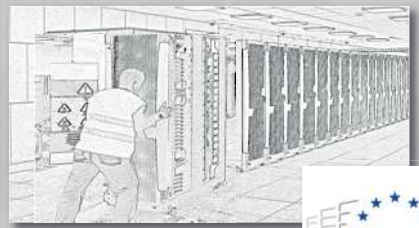

- Large cluster II
- 1000+ nodes
- €10-15mil.

EuroHPC
Joint Undertaking

+50%?

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

The small cluster II

- Atos BullSequana
- 189x compute nodes, 2x18 cores
  - AVX-512 instruction set, 192GB RAM
- 1x Fat node, 4x12 cores, 6TB RAM
- 8x GPU nodes, 2x12 cores, 192GB RAM
  - 4x Nvidia V100 per node, 16GB RAM
- Infiniband HDR, 200Gb/s link speed, 4 islands, fat tree topology
- Burst buffer accelerated SCRATCH 200TB, 30GB/s
- Small HOME, 25TB
- 14x1.6TB NVMe
  - accessible remotely on all nodes
- PBS-Pro scheduler
- 840TF Peak perfomance

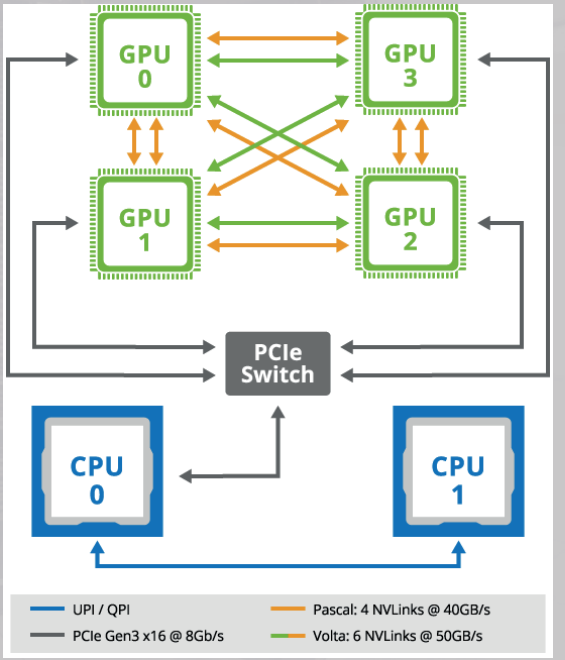

NVIDIA DGX-2

- 2x24 x86_64 cores
  - instruction set AVX-512
- architecture Volta V100 GPGPU
  - 16x2560 FP64 cores
  - 16x5120 FP32 cores
  - 16x640 tensor cores
- 1.5 TB RAM, 512GB HBM
- NVLINK network interconnecting GPGPU
  - 12x NVSwitch, throughput 2.4TB/s in bisection
- 8x 100Gb/s Infiniband
- NVMe SSD storage 30TB
- 130TF Peak!

IT4I infrastructure services

- Open access
  - Grant competition announced three times a year (February, June, October) for employees of research institutions, scientific and educational organizations
- Directors discretion
  - An application can be filed any time. Computing time is assigned irregularly based on an assessment by IT4Innovations
- National node of EU HPC infrastructures
  - Partnership for Advanced Computing in Europe (PRACE)
  - European Technology Platform for High (ETP4HPC)
  - EuroHPC Joint Undertaking (EuroHPC JU)
- Training and educational activities
  - More than 10 events annually, 63 events since 2013
  - 6 PRACE seasonal schools
  - PRACE Training Centre – 4 two days events in 2018-2019
- Users support

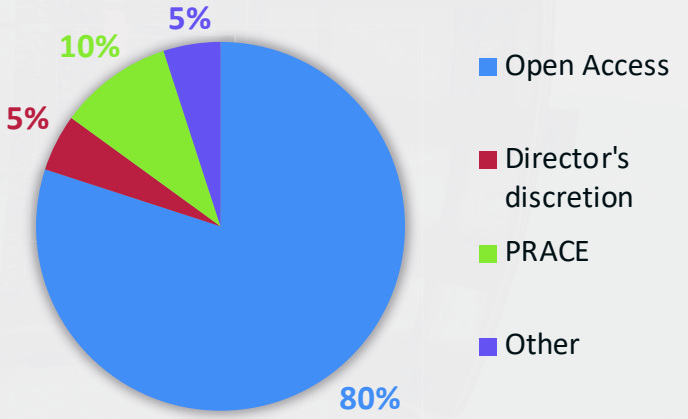

Open access

Call for proposals

- 3x per year with allocation period of 9 months
- Call is closed after 40 working days

Evaluation

- Peer review: Scientific, technical and socio-economic assessment, 2 external reviewers
- Applicant's history: citation of infrastructure, use of allocated resources
- In 20 working days, evaluation will be available
- 10 working days for rebuttal
- Access panel: Results announced in 5 working days

Changes suggested

- Multi-year projects: allocation period up to 4x9months, user's history, running project intermediate reports
- Evaluation based on applicant's project (H2020, CoE, CoC, VI, …)

Open access and utilization

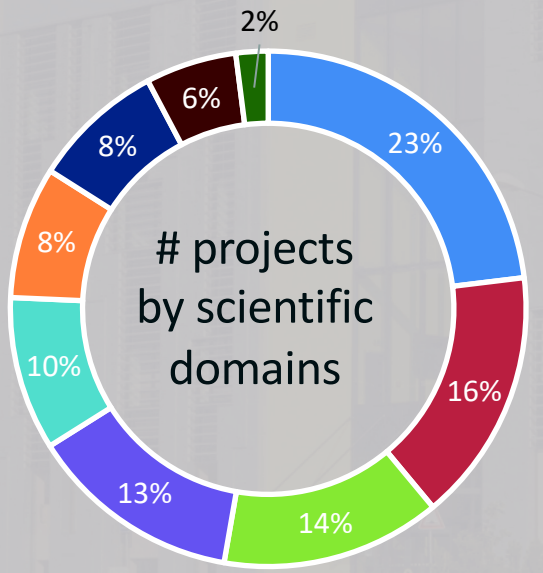

- Materials Science
- Bio Sciences
- Engineering
- Informatics
- Earth Sciences
- Training
- Applied Mathematics
- Plasma & Particle Physics
- Astro Sciences

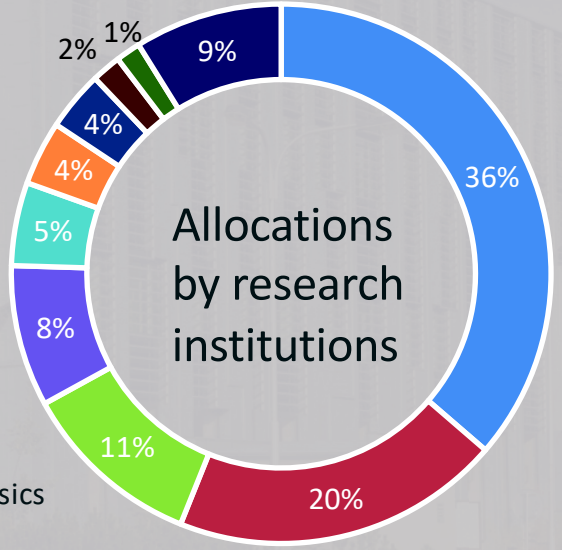

- IT4Innovations
- Akademie věd ČR
- CEITEC
- Univerzita Karlova
- Masarykova univerzita
- ČVUT
- VUT Brno
- Uppsala University
- KTH Royal Institute of Technology
- Others

In 2018, 171 584 400 corehours among 164 research projects allocated

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

Partnership for Advanced Computing in Europe

- **PRACE** is an international not-for-profit association under Belgian law, with its seat in Brussels.
- **PRACE** has 25 members and 2 observers.

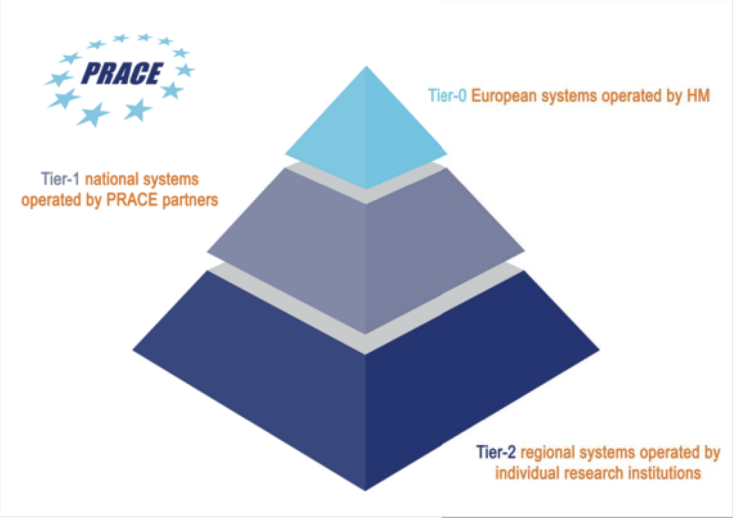

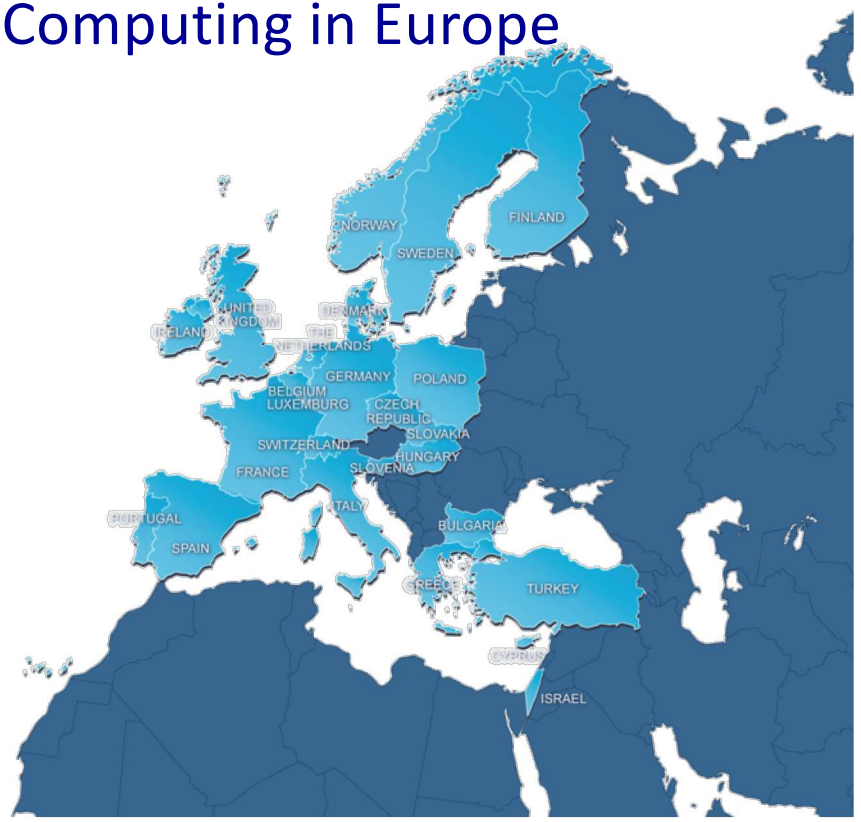

Development of HPC resources

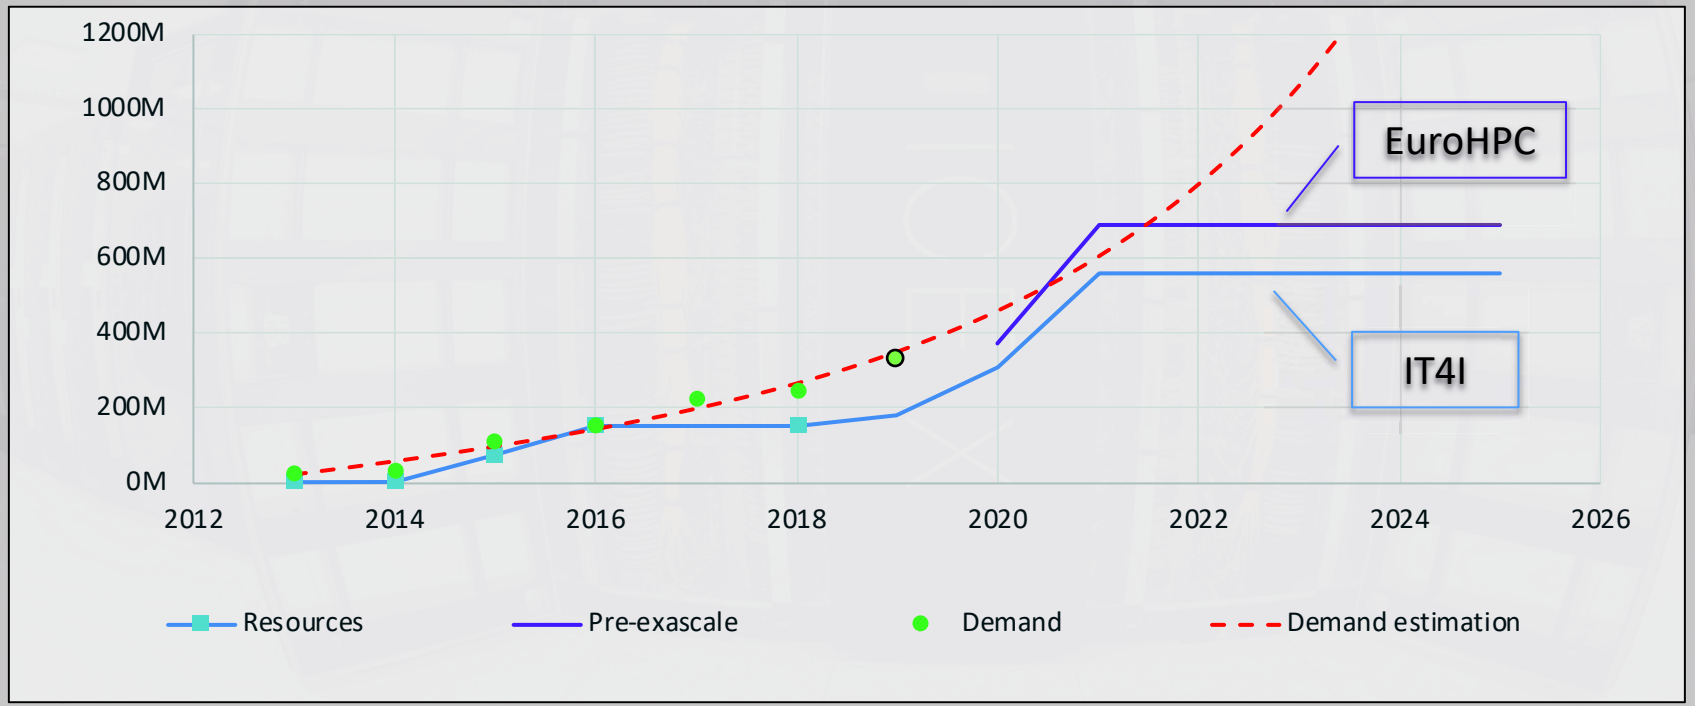

PRACE Tier-0 Systems

MARE NOSTRUM: IBM
BSC - Barcelona, Spain

#25 on Top500

PIZ DAINT: Cray XC40 / XC50
CSCS - Lugano, Switzerland

#5 on Top500

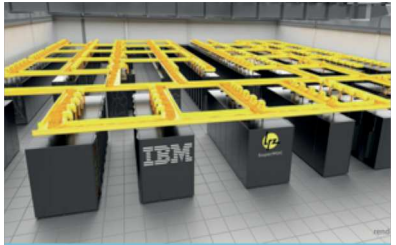

SUPERMUC: IBM
GAUSS/LRZ – Garching, Germany

#8 on Top500

JUWELL: Atos Sequanna X100,
GAUSS/FZJ, Germany

#26 on Top500

Joliot Curie: Atos Sequanna
X100, GENCI, France

#40 on Top500

MARCONI: Lenovo
CINECA - Bologna, Italy

#19 on Top500

HAZEL HEN: Cray
GAUSS/HLRS – Stuttgart, Germany

#30 on Top500

PRACE access models

Access Model for Tier-0

- Project access – large-scale, computationally intensive projects
- Preparatory access - Test / evaluation access, code scaling and optimization
- Proposals are evaluated in a single European peer review process governed by the PRACE Scientific Steering Committee

DECI (Distributed European Computing Initiative)

- Access to national Tier-1 systems
- Support by experts from Tier-1 centres

SHAPE (SME HPC Adoption programme in Europe)

- To help European SMEs overcome barriers to using HPC

Educational and Training activities

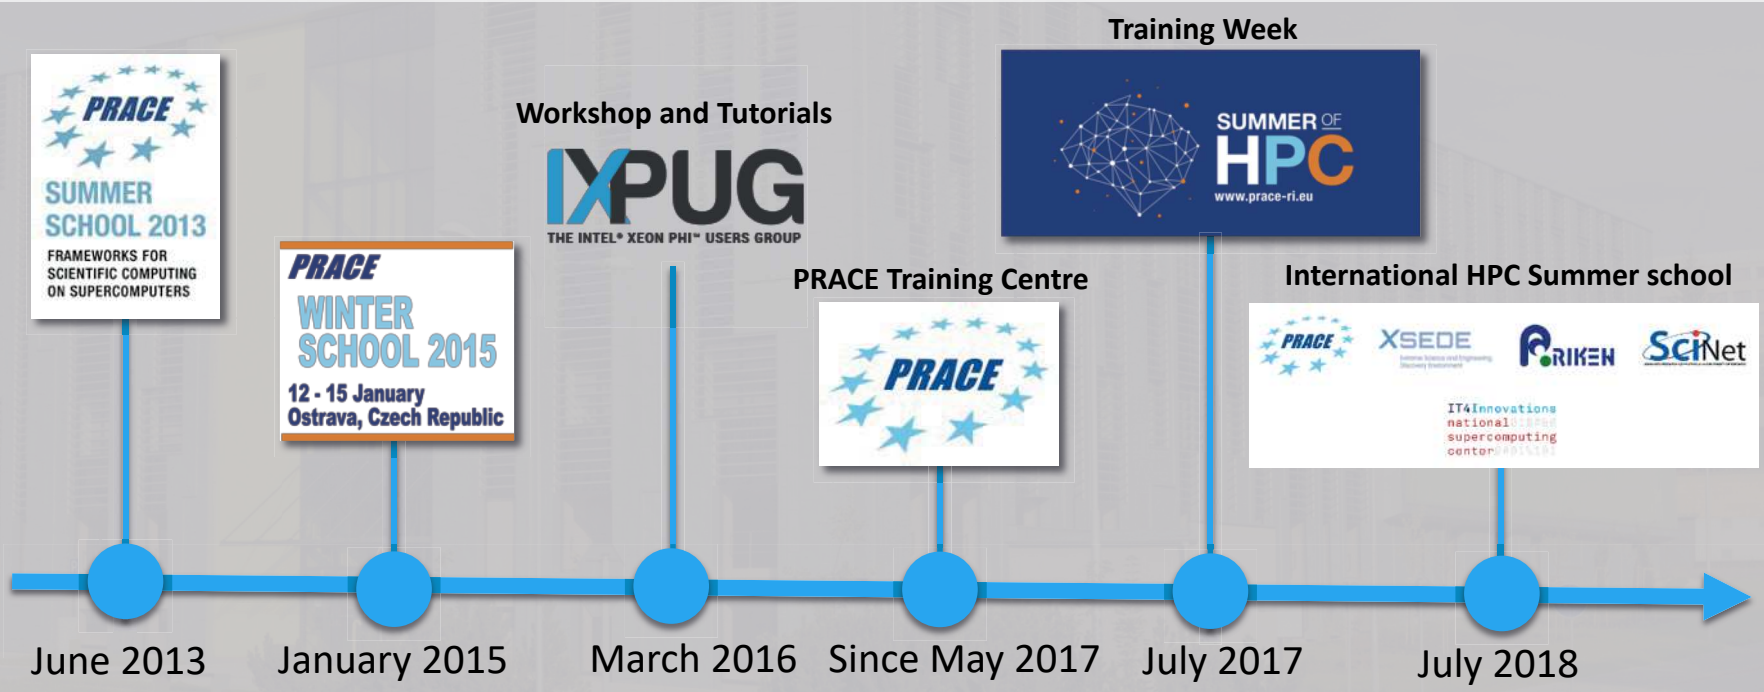

Since 2013 we are organizing approx. 10 courses per year in average. In 2017 we became PRACE Training Center.

European Commission

#DSM

Digital Single Market

POOLING RESOURCES FOR A EUROPEAN HIGH PERFORMANCE COMPUTING INFRASTRUCTURE

HIGH PERFORMANCE COMPUTING (HPC) – WHAT IS IT ABOUT?

Super powerful and efficient machines, able to process large amounts of data and perform calculations thousands of times faster than a normal computer.

10^{18}

The next step in HPC is exascale performance: a machine able to do 10^{18} (a billion of billion) calculations per second, a technology which is expected to be operational in 2021/2022.

HPC benefits

Citizens

Addressing major societal challenges of modern society (e.g. health, more efficient public services, cybersecurity, safer and greener transport)

Researchers and scientists

Underpinning innovation in almost all scientific disciplines

Deeper insights into unexplored systems of high complexity

Industry

+ innovative
+ efficient (resources and time)
- costly

HPC capabilities are used to solve and address scientific, industrial and societal challenges.

Health

- Development of personalised and precision medicine to provide individual and accurate patient treatment.
- Saving time and money on the development of new drugs from the initial idea phase to the final phase of reaching the market.

development of new drugs reduced from 17 years to several months

HPC can help save hundreds of millions up to billions of Euros

- Early detection of diseases and quicker diagnosis.

Climate change & weather forecast

- Europe paid severe weather damage costs between 1970 and 2012[1].

150.000 lives

€270 billion in economic damages

- With HPC technology, climate scientists will be able to predict the size and paths of storms and floods with accuracy.
- This will allow implementation of measures such as alerting or evacuating people, thus saving human lives.

Industry

- Reducing development time, minimising costs, optimising decision processes and producing higher quality goods and services.
- E.g. automotive industry will save time and money to develop new vehicle platforms, with improved environmental friendliness and passenger comfort and safety[2].

development reduced from 60 months to 24 months

HPC can help save up to €40 billions

Cybersecurity

- Enabling complex encryption technologies and better reactions to cyberattacks.
- Combined with Artificial Intelligence, HPC detects:
  - strange systems behaviour,
  - insider threats and electronic fraud,
  - very early cyberattack patterns (few hours instead of a few days)

This will allow automated and immediate actions even before a cyberattack happens.

Energy

- HPC provides critical tools for example in:
  - designing renewable energy parks
  - designing high-performance photovoltaic materials,
  - optimising turbines for electricity production.
- HPC expenditure in the energy sector is projected to grow by 5% in the next years.

Updated 04 June 2018

EuroHPC Joint Undertaking

WHY IS IT IMPORTANT TO JOIN FORCES AT EU LEVEL AND CO-INVEST IN WORLD-CLASS HPC INFRASTRUCTURES

At the moment, **EU industry provides** about **5% of HPC resources worldwide**, but **consumes one third** of them.

Industrial sectors that leverage HPC could **add up to 2-3% to Europe's GDP** in 2020 by improving their products and services.

97% of the industrial companies using HPC consider it **indispensable** for their ability to **innovate** and to **increase their competitiveness**.

In Europe, among 143 HPC projects analysed, **each €1 invested in HPC** has **generated** around **€870 in revenues** for businesses and **€69 in benefits**.

EU needs to have HPC world class systems. In June 2012, EU had 4 machines in the global top 10. Today the **fastest supercomputer in EU ranks 14** on the **global list** – About 12 times slower than the world's fastest machine.[3]

Compared to its competitors from USA, China or Japan, **Europe** is clearly **underinvesting in HPC** with a **funding gap of €500-750 million** per year[4].

HOW WILL EUROPE BECOME A WORLD LEADER IN HPC?

EuroHPC Declaration

The signatory **Member States** commit to **working together** with the Commission for **deploying** an **integrated world-class HPC infrastructure**.

One of the objectives is to have EU exascale supercomputers in the global top 3 ranking by 2022-2023.

The EuroHPC declaration has been launched on 23 March 2017 at the Digital Day in Rome.

#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.

EuroHPC Joint Undertaking

It is a legal and funding instrument to

- acquire and operate a world-class pre-exascale supercomputing infrastructure;
- develop exascale supercomputers based on competitive EU technology.

2019 – 2026

It is foreseen to **start operating from 2019 until 2026**.

It will be composed of public members (until now 13 Member States have signed up) and **private members** (representatives from HPC and Big Data stakeholders, including academia and industry).

It will get a budget of around **€1 billion**, of which 50% comes from the EU and 50% from the participating countries. Private entities will provide contributions in kind.

It will provide **financial support** (procurement or Research & Innovation grants) following **open and competitive calls**.

IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

EuroHPC Joint Undertaking

- CZ joined EuroHPC JU on January 24th, 2018
  - November 6th, 2018 – 1st Governing Board meeting
- Hosting EuroHPC petascale system
  - 3-4 petascale systems. To be installed in 2020
  - IT4I is going to apply for €5M support
- Membership in pre-exascale system consortium
  - 2 pre-exascale systems – world class systems €240M each
  - Finnish consorcium: FI, S, NO, DK, CH, BE, NL and CZ, PL(?)
- Research activities (EU funding up to €200M)
  - European processor initiative
  - Extreme scale technologies
  - Building HPC Competence centres + Skills + SMEs support
  - *Federating supercomputing resources, HPC applications*

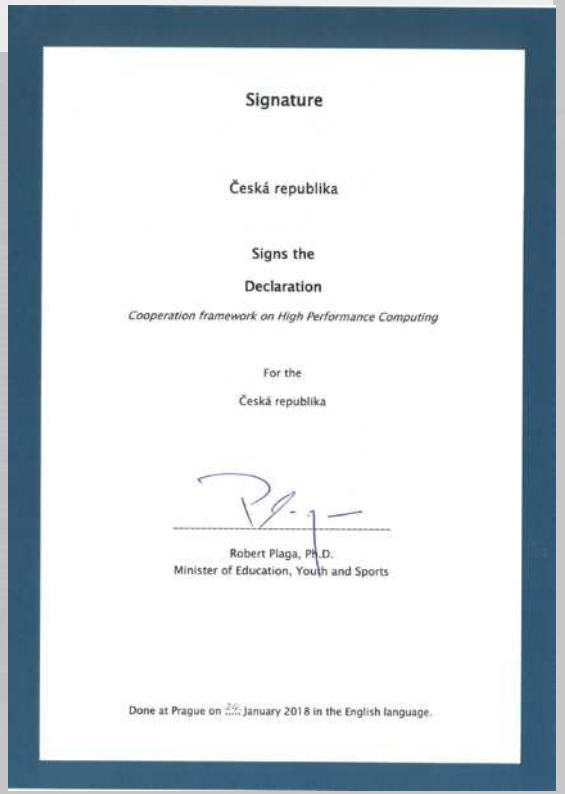

Signature

Česká republika

Signs the

Declaration

Cooperation framework on High Performance Computing

For the

Česká republika

Robert Plaga, Ph.D.
Minister of Education, Youth and Sports

Done at Prague on 24. January 2018 in the English language.

HPC and HPDA tools development

Research&Development

HyperLoom

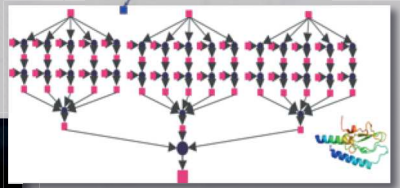

HPC as a Service

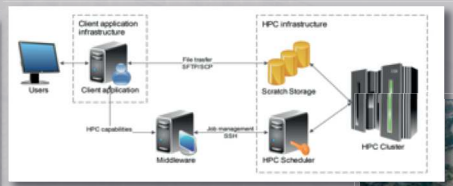

Society

HPC infrastructure

Energy efficiency

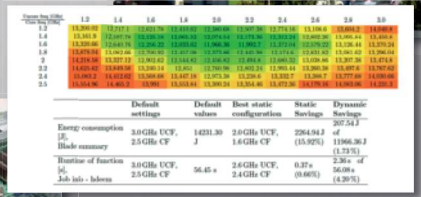

	Default settings	Default values	Best static configuration	Static Savings	Dynamic Savings
Energy consumption [J], Blade summary	2.0 GHz UCF, 2.5 GHz CF	14231.30 J	2.0 GHz UCF, 1.6 GHz CF	2264.94 J (15.92%)	207.54 J of 11966.36 J (1.73 %)
Runtime of function [s], Job info - hdeem	2.0 GHz UCF, 2.5 GHz CF	56.45 s	2.6 GHz UCF, 2.4 GHz CF	0.37 s (0.66%)	2.36 s of 56.08 s (4.20 %)

Scalable solvers

Industry

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

International projects

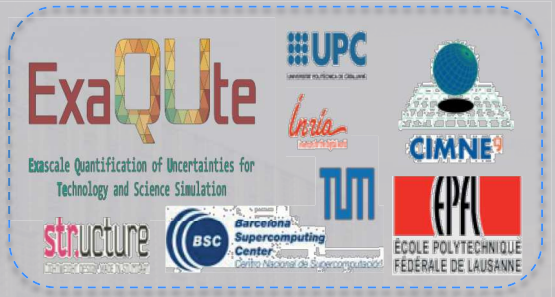

Industrial partners

Thank you for your attention!