High Performance Computing in the Czech Republic

Vít Vondrák IT4Innovations national supercomputing center



IT4Innovations national@13#20 supercomputing center@#81%101





Roadmap for Large Infrastructures of the Czech Republic 2016-2022

ICT/ e-infrastructures

| 10.6 | |
|------------------|-------|
| CERIT Scientific | Cloud |

IT4Innovations National Supercomputing Center

E-infrastructure CESNET

CESNET

high speed interconnect, grid computing, data archives and collaborative environment

IT4Innovations

Tier-1, national level supercomputers, HPC training, HPC application support

• CERIT-SC

flexibility in e-services, innovative use and experimental approach

Ministry of Education, Youth and Sports

[source: http://www.msmt.cz/vyzkum-a-vyvoj-2/cestovni-mapa-cr-velkych-infrastruktur-pro-vyzkum]

E-Infrastructures landscape in Czech Republic: Net, Storage, Grid, HPC



&#\$% IT4I @&01

IT4I infrastructure history

&#\$% **IT4I** @&01

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 (10³) Mega (10⁶) Giga (10⁹) Tera (10¹²) Peta (10¹⁵) Exa (10¹⁸)

Salomon in Top500.org

&#\$% IT4I @&01

| 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 IGENCI-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 Craving | 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/0, Power BQC 16C 1.60GHz, Custom IBM | 131,072 | 1,431.1 | 1,677.7 | 657 |
| 42 | Air Force Research Laboratory United States | Spirit - SGLICE 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 8510, Xeon E5-2680 8C 2.700GHz, Infiniband QDR Bull, Atos Group | 77,184 | 1,359.0 | 1,667.2 | 2,132 |

| List | Rank | System | Vendor | Cores | (TFlops) | (TFlops) | |
|---------|------|---|--------|--------|----------|----------|--|
| 11/2017 | 87 | SGI ICE X, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR, Intel Xeon Phi 7120P | HPE | 76,896 | 1,457.7 | 2,011.6 | |
| 06/2017 | 78 | SGI ICE X, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR, Intel Xeon Phi 7120P | HPE | 76,896 | 1,457.7 | 2,011.6 | |
| 11/2016 | 67 | SGI ICE X, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR, Intel Xeon Phi 7120P | HPE | 76,896 | 1,457.7 | 2,011.6 | |
| 06/2016 | 55 | SGI ICE X, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR, Intel Xeon Phi 7120P | HPE | 76,896 | 1,457.7 | 2,011.6 | |
| 11/2015 | 47 | SGI ICE X, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR, Intel Xeon Phi 7120P | HPE | 76,896 | 1,457.7 | 2,011.6 | |
| 06/2015 | 39 | SGI ICE X, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR, Intel Xeon Phi 7120P | HPE | 76,896 | 1,457.7 | 2,011.6 | |

Total

Rmax

Rpeak

IT4Innovations data storage

IT4I

Capacity 1.69PB Throughput 30GB/s

Anselm: Capacity 146TB Throughput 6GB/s

0.5PB Capacity Throughput 6GB/s

Anselm: 0.3PB Capacity Throughput 2GB/s Backup LTO-6 Capacity: 3PB R/W: 25TB/h

Data room infrastructure

&#\$% IT4I @&01

Dynamic rotating UPS 2x2,5MVA

OxyReduct fire prevention

Infrastructure services

Open access competition (80%)

- Grant competition announced three times a year (February, June, October) for employees of research institutions, scientific and educational organizations

• PRACE DECI access (10%)

 The Distributed European Computing Initiative (DECI) is designed for projects requiring access to resources not currently available in the PI's own country but where those projects do not require resources on the very largest (Tier-0) European Supercomputers

Directors discretion (5%)

 An application can be filed any time. Computing time is assigned irregularly based on an assessment by IT4Innovations

Training and educational activities

- More than 10 events annually, 56 events since 2013
- 6 PRACE seasonal schools
- PRACE Training Centre 4 two days events in 2018-2019
- User support

Utilization by institution

&#\$% IT4I @&01

Projects by institutions allocated within open access grant competitions in 2017

Utilizations by scientific domains

&#\$% IT4I @&01

Supercomputing in Science and Engineering (2017)

http://www.it4i.cz/info-publications

Marie Běhounková (Charles University): Computational planetology: Timing of water plume eruptions on Enceladus explained by interior viscosity structure

Michal Hradiš (Brno University of Technology): Convolutional Neural Networks for Image Restoration

Peter Huszár (Charles University): Urban canopy climate forcing on a regional scale modelled by the RegCM and WRF models

Pavel Hobza (The Czech Academy of Sciences): The SQM/COSMO Scoring Function Reliably Describes Noncovalent Interactions in Protein-Ligand Complexes

Pavel Jungwirth (The Czech Academy of Sciences): Conformational Transitions and Membrane Binding of the Neuronal Calcium Sensor Recoverin

Petr Vrchota (Czech Aerospace Research Centre): ACLA – Active Load Alleviation

Jaroslav Koča (Masaryk University): Enzymatic Reaction Mechanisms Investigated by Car-Parrinello ab initio dynamics

Jiří Jaroš (Brno University of Technology): Convergence testing of a k-space pseudospectral scheme for transcranial time-reversal focusing

Infrastructure utilization

IT4I

In 2017, 160 350 384 corehours allocated among 142 research projects allocated!

Building HPC community

13th Open access competition

- open 2/2/2018
- running until 19/2/2019
- allocation exceeded by 37%
 - (65 869 100, NaN, 48 000 000)
- 46 project requests

Call 6th Call 7th Call 8th Call 9th Call 10th Call 11th Call 12th

Infrastructure: plan 2018

&#\$% IT4I @&01

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

IT4Innovations National Supercomputing Centre – Path to Exascale

2018

Total investment: ~ € 3 mil. Performance: ~ 0.5 PFlops

- supercomputer Small cluster II
- 200 nodes
- up to 500TF
- GPU acceleration
- NVMe

Infrastructure: plan 2020

&#\$% IT4I @&01

Salomon

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

IT4Innovations National Supercomputing Centre – Path to Exascale

Total investment: ~ € 10 mil. Performance: ~ 10 PFlops 2020

- supercomputer Large cluster II
- 1000+ nodes
- HBM
- new technology???

Future trends – Green500 list

&#\$% **IT41** @&01

The exascale formula:

20 MW * 50 GF/W =

= 1000 PF

| Rank | Rank Top500 | System | Main proc. | Eff. [GF/W] | % Rpeak |
|------|----------------|-------------------|-------------|----------------|------------|
| 1 | 259 | Shoubu system B | PEZY-SC2 | 17.0 | 75 |
| 2 | 307 | Suiren2 | PEZY-SC2 | 16.8 | 73 |
| 3 | 276 | Sakura | PEZY-SC2 | 16.7 | 73 |
| 4 | 149 | DGX SaturnV Volta | Nvidia V100 | 15.1 | 58 |
| 5 | 4 | Gyoukou | PEZY-SC2 | 14.2 | 68 |
| 6 | 13 | TSUBAME3.0 | Nvidia P100 | 13.7 | 67 |
| 7 | 195 | AIST AI Cloud | Nvidia P100 | 12.7 | 45 |
| - | 4 (y 2013) | K computer | SPARC64 | 0.88 | 93 |
| 20 | 1 | Sunway TaihuLight | SW26010 | 6,1 | 74 |

The parallelization curse

| Processor | Parallelization | Rpeak [GF] | TDP [W] | Eff. [GF/W] |
|-----------------------|-----------------|------------|---------|-------------|
| Skylake (20c) | 1600 | 1600 | 150W | 10.6 |
| Nvidia P100 | 5376 1792*3 | 5300 | 300W | 17.6 |
| PEZY-SC2 | 4096 2048*2 | 4100 | 180W | 22.7 |
| Nvidia V100 | 7680 2560*3 | 7800 | 300W | 26.0 |
| Exascale Processor | (est. 15000) | 15000 | (300W) | 50 |

Exascale projects

HPC tools development

Scalable solvers for scientific and engineering problems

PRACE SHAPE

ING

Parallel framework for Engineering Application with Massively parallel sparse linear solver designed to take full advantage of todays most powerful petascale supercomputers the second second

20 billion DOF on up to 17 576 Compute Nodes (275k cores)

TT4T

Accessible HPC Platform for Executing Scientific Pipelines on HPC

High-End Application Execution

- HPC capabilities as-a-service
- Unified interface for different schedulers
- Authentication and authorization
- Monitoring and reporting

HyperLoom

- Defining and executing scientific pipelines
- Custom scheduling algorithm
- Execution of millions of interconnected tasks on hundreds of computational nodes
- Python client for pipelines definition and execution

Educational and Training activities IT4I **Training Week** Workshop and Tutorials SUMMER OF SUMMER PUG **SCHOOL 2013** www.prace-ri.eu FRAMEWORKS FOR SCIENTIFIC COMPUTING ON SUPERCOMPUTERS **PRACE Training Centre** PRACE International HPC Summer school XSEDE CRIKEN SCINET SCHOOL 2015 PRACE 12 - 15 January IT4Innovations **Ostrava, Czech Republic** national01 supercomputing center® March 2016 Since May 2017 June 2013 January 2015 July 2017 July 2018

Training courses 2017-2018

0&01

- 10.05.2018 11.05.2018 PETSc Basic & Advanced Tutorial (PRACE Training Course) Václav Hapla (ETH Zürich, Switzerland)
- 23.04.2018 27.04.2018 VI-HPS Tuning Workshop (PATC course, LRZ, Germany) developers of the VI-HPS tools
- 19.04.2018 Aktuální trendy v programování GPU Ivan Šimeček (GPU Education Center při FIT ČVUT)
- 22.03.2018 23.03.2018 Parallel I/O & Libraries (PRACE training course) Nicole Audiffren (CINES, France), Sebastian Lührs (JSC, Germany)
- 01.03.2018 02.03.2018 Intel Xeon Phi programming (PTC course) Georg Zitzlsberger, Martin Golasowski, Michal Merta (IT4Innovations)
- 24.01.2018 High Performance Distributed Deep Learning Dhabaleswar K. Panda, Hari Subramoni, Ammar Ahmad Awan (The Ohio State University)
- 23.01.2018 InfiniBand, Omni-Path, and High-Speed Ethernet: Advanced Features, Challenges in Designing, HEC Systems and Usage Dhabaleswar K. Panda, Hari Subramoni (The Ohio State University)
- 23.01.2018 InfiniBand, Omni-Path, and High-Speed Ethernet for Dummies Dhabaleswar K. Panda, Hari Subramoni (The Ohio State University)
- 14.12.2017 15.12.2017 CFD simulations using OpenFOAM (PTC course) Tomáš Brzobohatý (IT4Innovations)
- 27.11.2017 28.11.2017 Productivity tools for High Performance Computing (PTC course) Branislav Jansík, David Hrbáč, Josef Hrabal, Lukáš Krupčík, Lubomír Prda, Roman Slíva (IT4Innovations)
- 04.10.2017 05.10.2017 Creating Robust Software for Better Scientific Outcome Anshu Dubey, Rinku Gupta (Argonne National Laboratory, USA)
- 14.06.2017 What you need to know about performance analysis using Intel tools Georg Zitzlsberger (IT4Innovations)
- 27.04.2017 Efficient HPC Development and Production with Allinea Tools Florent Lebeau (Allinea, UK)
- 02.03.2017 03.03.2017 GASPI Global Address Space Programming Interface Christian Simmendinger (T-Systems Solutions for Research GmbH)
- 07.02.2017 08.02.2017 Intel MIC Programming Workshop Volker Weinberg, Momme Allalen (Leibniz Supercomputing Centre), Branislav Jansík (IT4Innovations)
- 02.02.2017 03.02.2017 Parallel Linear Algebra Mathieu Faverge, Florent Pruvost, Gilles Marait (Bordeaux INP), Zdenek Strakos (Charles University)

Participation in EuroHPC JU

The In-kind contribution working group

Identify concrete items for in-kind contributions

тта

- Establish best practices and methods
- Treat both pilar 1 and pilar 2
- ČR rep. by Branislav Jansik (IT4I)
- The User requirements working group
- Report on user requirements
- Draft machine architecture requirements
- Benchmarking suites and criteria
- ČR rep. by Branislav Jansik (IT4I)

The Hosting and Procurement working group

- Hosting selection criteria
- Hosting agreements
- Procurement process and role of hosting centers
- ČR rep. by Filip Stanek (IT4I)

Contribution to EuroHPC JU

Research activities

- Development of algorithms and parallel programming models
- Software tools, platforms and libraries
- Code enabling and optimization for future exascale systems
- HPC infrastructure operation and provisioning
 - EuroHPC Petascale system joint procurement with planed upgrade of Large cluster II

TTAI

Procurement and hosting conditions and requirements to be defined

Training activities

- Courses on use of HPC and HPDA infrastructures
- Basic, intermediate and advanced courses on parallel programming and use of software tools
- User support and assistance

