




IT4INNOVATIONS SUPERCOMPUTING INFRASTRUCTURE AND ITS EFFICIENT USE IN SCIENCE AND INDUSTRY

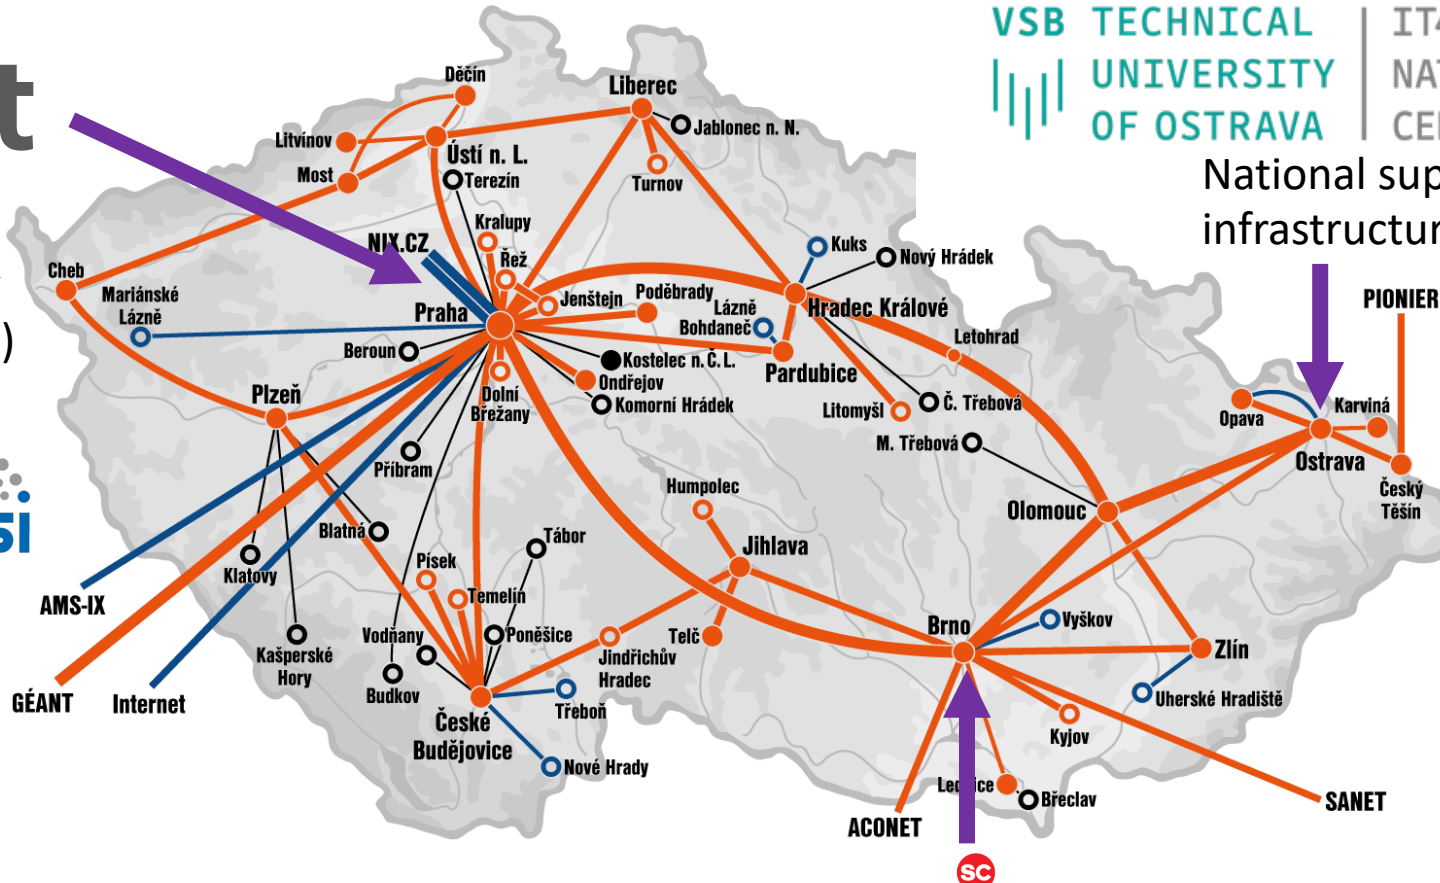
TOMÁŠ KOZUBEK

2023

E-INFRA CZ



NREN/Connectivity
MetaCentrum (NGI)
Data storage



Access to supercomputers and related services



National supercomputing infrastructure



EuroHPC
Joint Undertaking




LARGE RESEARCH
INFRASTRUCTURES



Largest node of
MetaCentrum



WHAT IS A SUPERCOMPUTER?



- a **very powerful computer**, most recently in the form of a **computer cluster**, i.e. a **connection of many ordinary computers together**, usually by a special **high-speed network**

Compute nodes




Data storage



Network (interconnect)



COMPUTATIONAL POWER




FLOP = Floating point operation

Computational complexity - number of floating point operations FLOPs

Computing power - number of floating point operations per second FLOPS (FLOP/s)

Value	Name	Shortcut	Exp.
1 000	kiloFLOPS	kFLOPS	10^3
1 000 000	megaFLOPS	MFLOPS	10^6
1 000 000 000	gigaFLOPS	GFLOPS	10^9
1 000 000 000 000	teraFLOPS	TFLOPS	10^{12}
1 000 000 000 000 000	petaFLOPS	PFLOPS	10^{15}
1 000 000 000 000 000 000	exaFLOPS	EFLOPS	10^{18}
1 000 000 000 000 000 000 000	zettaFLOPS	ZFLOPS	10^{21}

LINPACK BENCHMARK



- Theoretical performance - $R_{peak} = \#cores \times frequency \times \#FLOPs/cycle$
- LINPACK performance – based on the solution of the system of linear equations - R_{max}

```
Thu Sep 26 23:39:14 CEST 2019
Sample data file lininput.

Current date/time: Thu Sep 26 23:39:14 2019

CPU frequency:      3.789 GHz
Number of CPUs: 1
Number of cores: 4
Number of threads: 4

Parameters are set to:

Number of tests: 9
Number of equations to solve (problem size) : 15000 14000 13000 12000 11000 10000 8000 6000 1000
Leading dimension of array                   : 15000 14008 13000 12008 11000 10008 8008 6008 1000
Number of trials to run                      : 1 2 2 2 2 2 2 3 4
Data alignment value (in Kbytes)            : 4 4 4 4 4 4 4 4 4

Maximum memory requested that can be used=1800304096, at the size=15000


===== Timing linear equation system solver =====

Size  LDA  Align. Time(s)  GFlops  Residual  Residual(norm)  Check
15000 15000 4 15.121 148.8321 2.206862e-10 3.475844e-02 pass
14000 14008 4 12.115 151.0290 1.858561e-10 3.356364e-02 pass
14000 14008 4 12.337 148.3099 1.858561e-10 3.356364e-02 pass
13000 13000 4 9.813 149.2988 1.490012e-10 3.117738e-02 pass
13000 13000 4 9.818 149.2154 1.490012e-10 3.117738e-02 pass
12000 12008 4 7.624 151.1339 1.267737e-10 3.111642e-02 pass
12000 12008 4 7.637 150.8754 1.267737e-10 3.111642e-02 pass
11000 11000 4 5.762 154.0339 1.183327e-10 3.452256e-02 pass
11000 11000 4 5.720 155.1702 1.183326e-10 3.452256e-02 pass
10000 10008 4 4.124 161.6949 1.034086e-10 3.646293e-02 pass
10000 10008 4 4.188 159.2393 1.034086e-10 3.646293e-02 pass
8000 8008 4 2.057 165.9998 6.374379e-11 3.506461e-02 pass
8000 8008 4 2.062 165.5858 6.374379e-11 3.506461e-02 pass
6000 6008 4 0.888 162.2996 3.737388e-11 3.624465e-02 pass
6000 6008 4 0.908 158.6364 3.737388e-11 3.624465e-02 pass
6000 6008 4 0.886 162.5910 3.737388e-11 3.624465e-02 pass
1000 1000 4 0.005 136.1518 1.502479e-12 5.123838e-02 pass
1000 1000 4 0.005 135.1089 1.502479e-12 5.123838e-02 pass
1000 1000 4 0.005 132.4607 1.502479e-12 5.123838e-02 pass
1000 1000 4 0.005 134.2681 1.502479e-12 5.123838e-02 pass
```

$R_{max} = 166GFLOPS$
 $R_{max}/R_{peak} = 68\%$



THE MOST POWERFUL SUPERCOMPUTERS



- 2x a year ranking of the most powerful systems of the world
- www.top500.org




Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	Frontier - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,730,112	1,102.00	1,685.65	21,100

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	Frontier - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,730,112	1,102.00	1,685.65	21,100
2	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	29,899
3	LUMI - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,220,288	309.10	428.70	6,016
4	Leonardo - BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, Atos EuroHPC/CINECA Italy	1,463,616	174.70	255.75	5,610
5	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148.60	200.79	10,096



IT4INNOVATIONS INFRASTRUCTURE




Anselm (0,1PF)




IT4I building



Salomon (2,0PF, #40)



Barbora (0,8PF)



VaVR labs



Karolina (15,7PF , #69)



LUMI (550+PF, #3, Kajaani)




LUMI-Q



1PF = 10^{15} floating point operations per second

IT4INNOVATIONS SUPERCOMPUTERS





Anselm (0,1PF)

IT4I building

Salomon (2,0PF)

Barbora (0,8PF)

Karolina (15,7PF)



July 2013

July 2014

July 2015


October 2019

August 2021


NAMES

- Salomon – Salomon Mayer von Rothschild, the founder of Vítkovice Mining and Iron Corporation
- Anselm – son of Salomon, decommissioned in 2021 and moved to the Science Center in Ostrava-Vítkovice
- Caroline – wife of Salomon (in Czech: Karolina)
- Barbora – Patron Saint of Miners, name of the mine

EUROHPC JOINT UNDERTAKING



- 32 European countries
- Czechia is a member from January 2018



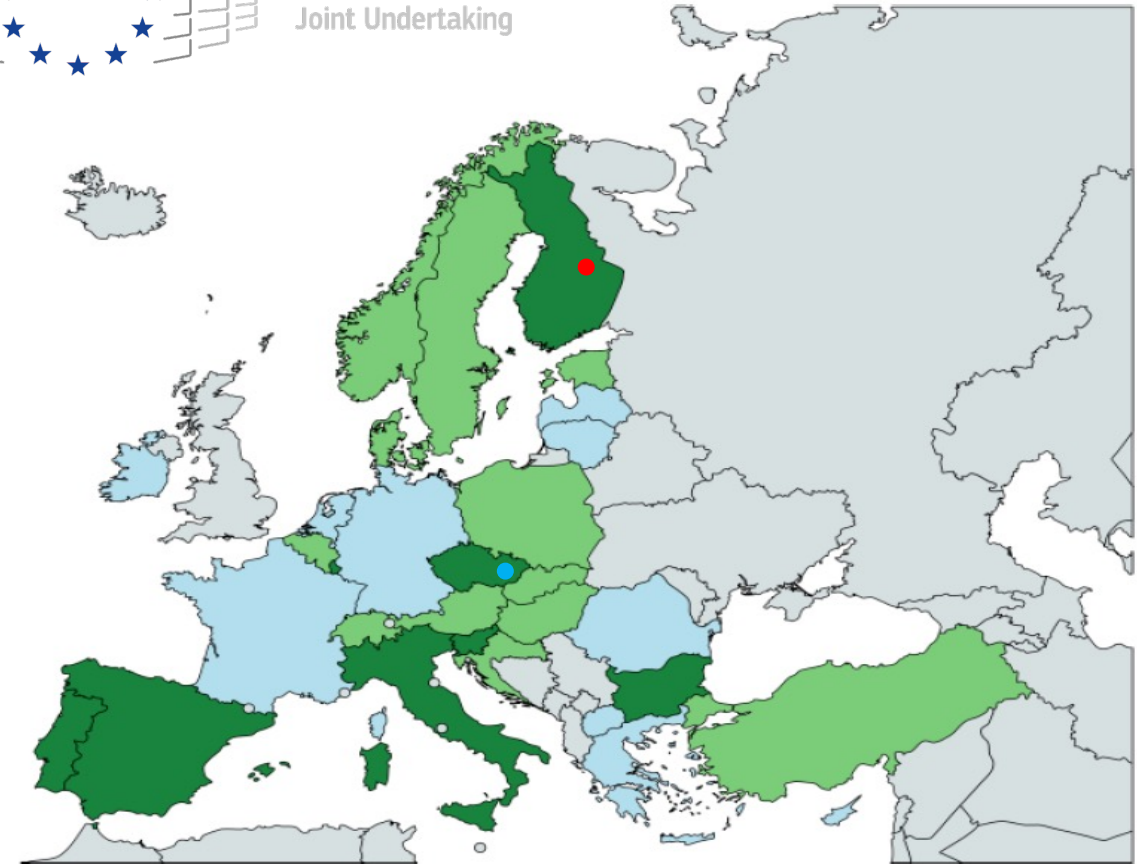
EuroHPC
Joint Undertaking

Build European HPC infrastructure


- 5 petascale supercomputers
 - **Karolina (CZ)**
 - Meluxina (LU)
 - Deucalion (PT, ES)
 - Vega (SI)
 - PetaSC (BG)
- 3 pre-exascale supercomputers
 - **LUMI consortium (FI, SE, NO, DK, BE, PL, CH, CZ, EE, IS)**
 - BSC (ES, ES, HR, PT, TR)
 - Leonardo (IT, SI, SK, AT, HU)

Support Research and Innovation in HPC

- HPC Centers of Competence (EuroCC)
- EuroHPC projects: LIGATE, ACROSS, IO-SEA, LIGATE, SCALABLE, EUmaster4HPC
- H2020: LEXIS, POP2, EVEREST, OPENQKD, EXAQUTE, TETRAMAX, DICE, CLOUDIFACTURING, EUPEX



KAROLINA SUPERCOMPUTER



Peak performance: 15.7 Pflop/s

Period of operation: 2021–2025

Total acquisition cost: 15M€ (35% EuroHPC JU)

Total operating cost: 13,9M€

- **Universal partition:** 720 nodes, 92160 AMD CPU cores
- **Accelerated partition:** 72 GPU nodes, 576 A100 GPUs
- **Data analytics partition:** 1 node, 24 TB RAM
- **Cloud partition:** 36 nodes, 4608 cores, 9 TB RAM

35 % of resources available for EuroHPC member states


TOP500 ranking – GPU partition:

- #69 (worldwide), #19 (Europe)



Green500 ranking – GPU partition:



- #8 (worldwide), #3 (Europe)



EuroHPC
Joint Undertaking




LUMI supercomputer, CSC-FI



- | Manufactured by **Hewlett Packard Enterprise**
- | Total budget: 207.1M€ (**50 % EuroHPC**)
- | #3 in TOP500 and GREEN500
- | Period of operation: 2021-2026
- | Consortium: **FI (CSC in Kajani)**, SE, NO, DK, BE, PL, CH, **CZ**, EE, IS
- | Ca 3,5% of the resources available to the Czech users through IT4I Open Access calls

1 system
550
Pflop/s
 Peak Performance


Computing power
 equivalent to
1 500 000
 Modern laptop computers




Size of a tennis court

Modern platform for
High-performance
computing,
Artificial intelligence,
Data analytics
Based on GPU technology


Visualization and virtual reality lab



- Barco's top-of-the-line 3D laser cinema projector with high 4K resolution,
- 100 Gb/s Ethernet line to transmit data and images directly from the IT4Innovations supercomputers,
- cooperation of up to 4 people in virtual reality.




TRAINING



- More than 15 courses annually in the fields of HPC, HPDA, AI, and QC
- Customised courses for clients
- From elementary to advanced level
- Certified instructors
- PRACE Training Centre
- Workshops and conferences
- Courses examples
 - | Fundamentals of Deep Learning for Multi-GPUs
 - | Introduction to HPC
 - | Introduction to Atos QLM and how to use it to run your first quantum circuit
 - | Data science with R and Python (PTC course)
 - | High-Performance CFD using OpenFOAM
 - | Introduction to Performance and Energy Efficiency Analysis (PTC course)
 - | Access to Karolina



HUMAN RESOURCES



Employees of IT4Innovations by divisions in full time equivalent (FTE), ca **150 FTE in total**:



21%

**Management and
Administration**



65%

**Research and
Development**




14%

**Supercomputing
Services**

Research laboratory	FTE
Parallel algorithms research (engineering) lab	16
Advanced data analysis and simulations lab	40
Modelling for nanotechnologies lab	18
Big data analysis lab	4
Infrastructure research lab	16

JOINT PROJECT ACTIVITIES



Research & Development



Support to industry



Infrastructure development+




Education




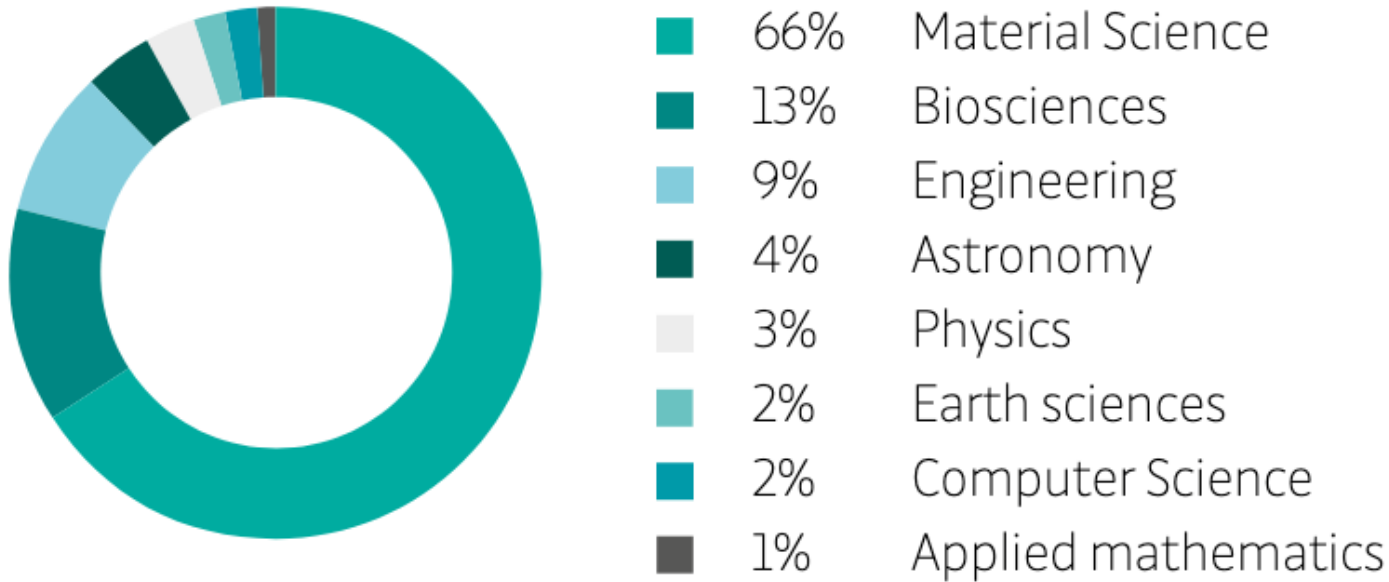
Partnership



WHAT ARE THE SCIENTISTS COMPUTING?



Computational resources allocated within the Open Access Grant Competitions in 2021 by scientific discipline




Optimization design of functional materials in a new type of lithium based battery


Dominik Legut (IT4I)




Parallel solvers for groundwater flow
Jakub Šístek (CAS)



Ab initio molecular dynamics study of the solvated electron in liquid ammonia
Pavel Jungwirth (IOCB)



Convection and climate simulations
Michal Belda (CUNI)




Simulation and planning of ultrasound surgical procedures
Jiří Jaroš (BUT)



New drug design
Pavel Hobza (IOCB)

WHO ARE OUR USERS?




Computational resources allocated within the Open Access Grant Competitions in 2021 by institution



31%	VSB - Technical University of Ostrava
26%	The Czech Academy of Sciences
8%	Charles University
8%	The Czech Technical University in Prague
5%	Masaryk University
4%	CEITEC
4%	Palacký University
4%	The University of Chemistry and Technology in Prague
4%	The University of Ostrava
3%	The University of Hradec Králové
2%	Brno University of Technology
1%	Others (Tomas Bata University in Zlín, the University of South Bohemia in České Budějovice, the Czech Aerospace Research Centre)




HPC, HPDA, and AI infrastructure & services
The most powerful public supercomputing infrastructure in the Czech Republic
VSB-TUO & MSIC



Infrastructure provider ★ HPC adoption programme ★ Collaborative research ★ Training & education

Advanced computing
and simulations



Virtual prototyping
Digital twins



Big data
Machine learning



Visualisation
Virtual reality



EuroHPC
Joint Undertaking

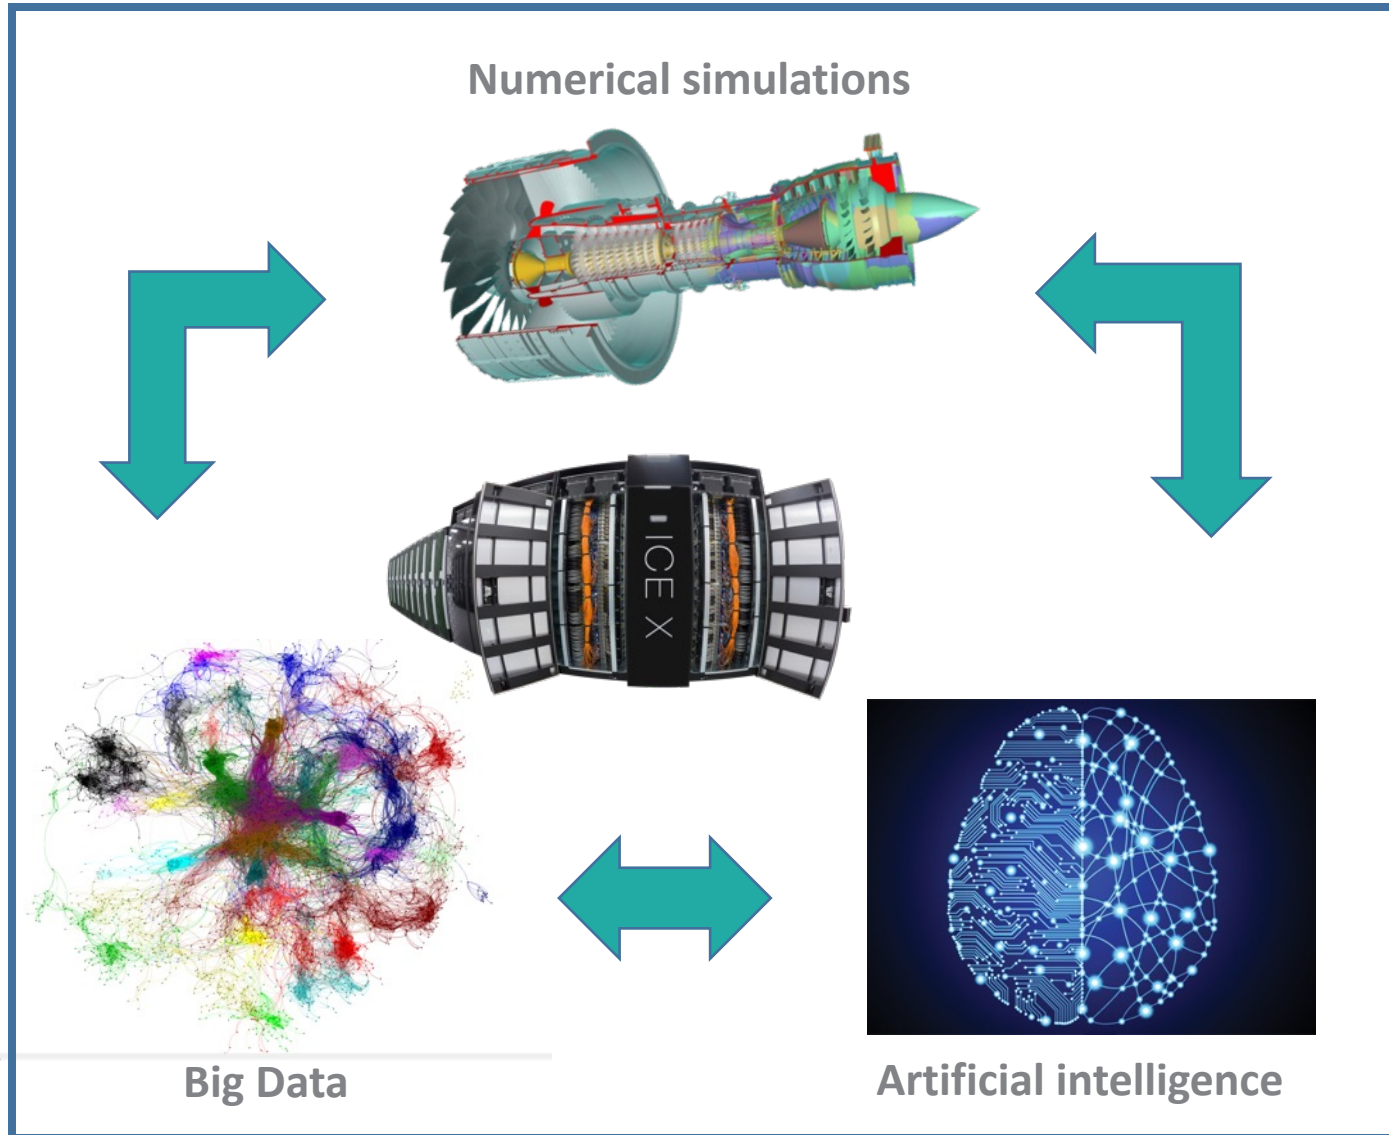



National Competence Centre in HPC from 2020

VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

RESEARCH ACTIVITIES



Research & Development



Society




Industry



Digital Twin of Siemens Electric Motor

Complex nonlinear multiphysical problem

- Electric fields
- Electromagnetism
- Heat transfer
 - heat generated by magnetism
 - cooling system
- Structural Mechanics
 - structural integrity
 - vibration from motion
 - high speed motors
 - influenced by electromagnetism
- Active cooling system
 - fluid flow
- Acoustic
 - generated by fluid flow
 - generated by electromagnetism
 - generated by vibrations





**Further acceleration:
AI based surrogate models**

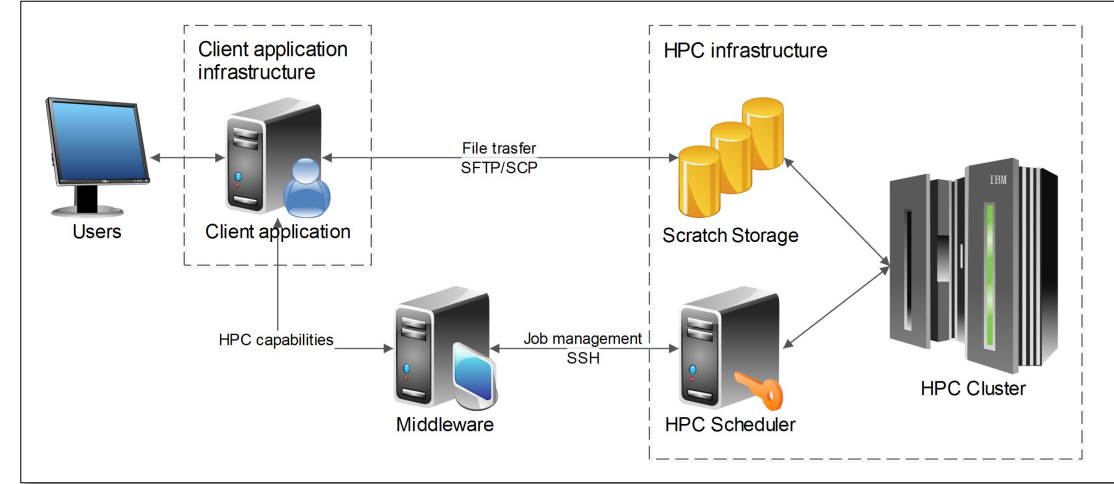

One complex transient simulation

- Salomon: 1200 CPU cores **by open source code**, 35 hours
- Workstation: 256GB RAM, 30 cores, more than 2 months!


Simulator for Generating Training Data for the Obstacle Detection System on a Railway Track



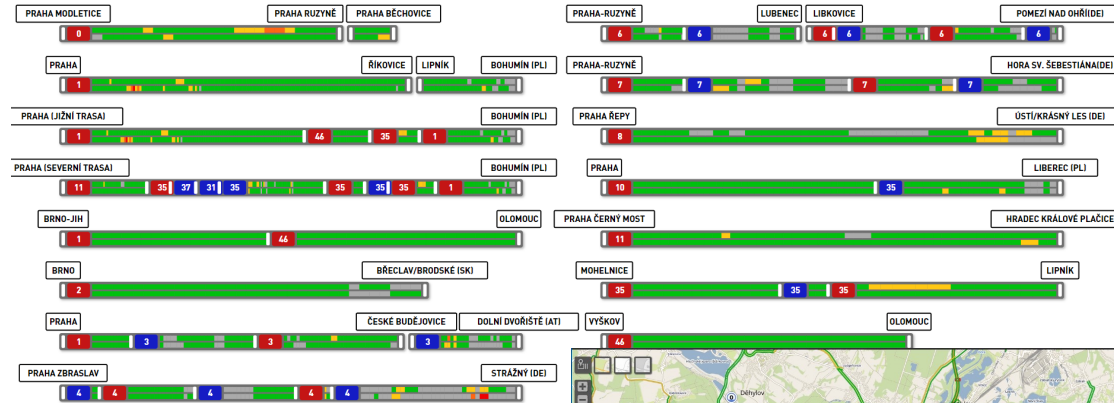
SMART PLATFORMS: VIARODOS & FLOREON+



HEAPPE, HYPERTOOLS

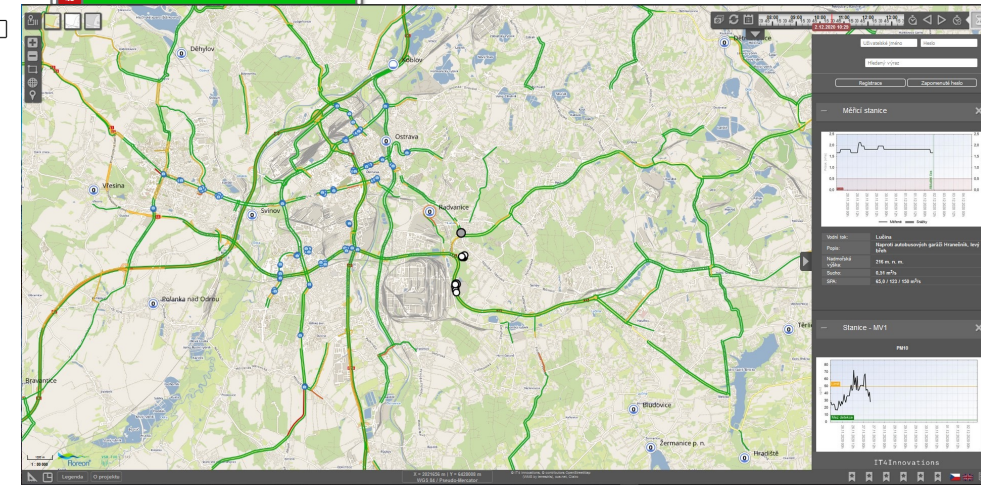


National Traffic Information Centre




Average amount of processed data

- Per working day:
15GB (8 million records / cca 25k time series)
- Per weekend day:
3GB (3 million records / cca 25k time series)



WEB IT4I.CZ, DOCS.IT4I.CZ



VSb TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

EUROCC E-INFRA.CZ EXTRANET | CZ

ABOUT INFRASTRUCTURE RESEARCH INDUSTRY COOPERATION FOR USERS EDUCATION EVENTS Q

WHAT DO OUR SUPERCOMPUTERS SOLVE?

ASTROPHYSICS

ENGINEERING

EARTH SCIENCES

CHEMISTRY

LIFE SCIENCES

MATERIAL SCIENCES

EVENTS

Every year we organize dozens of events for professionals and the public.

08/06 2022
Fundamentals of Deep Learning for Multi-GPUs
→ FIND MORE

14/06 2022
Introduction to HPC
→ FIND MORE

17/06 2022
EOSC Roadshow Ostrava
→ FIND MORE

IT4Innovations Documentation

Search

General Storage Clusters Software

Documentation

General

- [Introduction](#)
- Get Access >
- Get Project >
- Manage Your Profile >
- Access Services >
- Run Jobs >
- Technical Information >
- Satisfaction and Feedback
- PRACE
- DICE
- Support

Welcome to the IT4Innovations documentation. The IT4Innovations National Supercomputing Center operates the [Karolina](#) and [Barbora](#) supercomputers. The supercomputers are available to the academic community within the Czech Republic and Europe, and the industrial community worldwide. The purpose of these pages is to provide comprehensive documentation of the hardware, software, and usage of the computers.

How to Read the Documentation

1. Select the subject of interest from the left column or use the Search tool in the upper right corner.
2. Scan for all the notes and reminders on the page.
3. If more information is needed, read the details and **look for examples** illustrating the concepts.



Required Proficiency

Note

Basic proficiency in Linux environments is required.

In order to use the system for your calculations, you need basic proficiency in Linux environments. To gain this proficiency, we recommend you read the [introduction to Linux](#) operating system environments, and install a Linux distribution on your personal computer. For example, the [CentOS](#) distribution is similar to systems on the clusters at IT4Innovations and it is easy to install and use, but any Linux distribution would do.

GUIDED TOUR THROUGH THE INFRASTRUCTURE





prof. Tomáš Kozubek
tomas.kozubek@vsb.cz

IT4Innovations National Supercomputing Center
VSB – Technical University of Ostrava
17. listopadu 2172/15
708 00 Ostrava-Poruba, Czech Republic
www.it4i.cz

VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER