COMPUTING INFRASTRUCTURE



OUR SUPERCOMPUTERS SUPPORT EUROPEAN SCIENCE, INDUSTRY, AND SOCIETY

IT4Innovations National Supercomputing Center at VSB - Technical University of Ostrava is a leading research, development, and innovation centre active in the fields of High-Performance Computing (HPC), Data Analysis (HPDA), Quantum Computing (QC), and Artificial Intelligence (AI) and their application to other scientific fields, industry, and society. IT4Innovations operates the most powerful supercomputing systems in the Czech Republic, which are provided to Czech and foreign research teams from both academia and industry. Together with the CESNET and CERIT-SC institutions. IT4Innovations constitutes e-INFRA CZ, a strategic research infrastructure of the Czech Republic.

IT4Innovations currently operates three supercomputers — Barbora, NVIDIA DGX-2, a specialized system for AI calculations, and an Euro-HPC petascale system called Karolina with a theoretical peak performance of about 15.7 PFlop/s.

Czech research communities also have access to the LUMI supercomputer thanks to IT4Innovations' membership in the LUMI consortium. LUMI, with its theoretical performance of 580+ PFlop/s is the most powerful supercomputer in Europe and is located in Kajaani, Finland. IT4Innovations also participates in its operation.

In 2023, IT4Innovations signed a hosting agreement with the EuroHPC JU as leader of the LUMI-Q consortium. Under this agreement, one of the six EuroHPC quantum computers in Europe will be installed and operated at IT4Innovations.

Computational Resources Allocation

The computational resources of IT4Innovations are dedicated to solving problems from research and development done by academic and research institutions. Part of the capacity is dedicated to development of collaboration between academia and industry, or for the industry itself.

Open Access - this category of access aims to allocate computational resources to the research community in the Czech Republic based on scientific excellence, computational competence and readiness, and anticipated benefits to society and the economy. Computing resources are distributed in an open access model, which aims to make HPC resources available for non-economic

research and development activities. The largest share of the computational resources is distributed within the framework of Open Access Grant Competitions.

Access for Thematic HPC Resource Utilisation - allocations for socially important tasks, teaching and educational activities, commercial activities (contract research), and in-house infrastructure research. Submissions at any time.

Access to IT4Innovations computational resources can also be gained through EuroHPC JU Grant Competitions announced by the EuroHPC JU.

IT4I is a proud member of













VSB TECHNICAL | ITAINNOVATIONS |||| UNIVERSITY | NATIONAL SUPERCOMPUTING OF OSTRAVA | CENTER

0

.

Technical Specifications

	NVIDIA DGX-2	BARBORA	KAROLINA	LUMI
Put into operation	Spring 2019	Autumn 2019	Summer 2021	Autumn 2022
Theoretical peak performance	130 TFlop/s	849 TFlop/s	15.7 PFlop/s	580+ PFlop/s
Operating system	CentOS 7	RHEL 8	CentOS 7	HPE Cray OS
Compute nodes	1	201	831	5,042
Types of compute nodes		192 CPU nodes 2x Intel Cascade Lake 6240, 18-core, 2.6 GHz, 192 GB RAM	756 CPU nodes 2x AMD EPYC 7h12, 64-core, 2.6 GHz, 256 GB RAM	2,048 CPU nodes 2x AMD EPYC 7763, 64-core, 2.45 GHz, 256-1024 GB RAM
	1 GPU node 2x Intel Xeon Platinum 8168, 24-core, 2.7 GHz, 1.5 TB RAM, 16x NVIDIA Tesla V100, 32 GB HBM2	8 GPU nodes 2x Intel Skylake 6126, 12-core, 2.6 GHz, 192 GB RAM, 4x NVIDIA Tesla V100, 16 GB HBM2	72 GPU nodes 2x AMD EPYC 7763, 64-core, 2.45 GHz, 1 TB RAM, 8x NVIDIA A100, 40 GB HBM2	2,978 GPU nodes 1x AMD EPYC 7A53, 64-core, 2.45 GHz, 512 GB RAM, 4x AMD Instinct MI250X GPUs, 128 GB HBM2e
		1 data analytics node 8x Intel Xeon 8153, 16-core, 2.0 GHz, 6 TB RAM	1 data analytics node 32x Intel Xeon-SC 8628, 24-core, 2.9 GHz, 24 TB RAM	8 data analytics nodes 2x AMD EPYC 7742, 64-core, 2.25 GHz, 4 TB RAM
			2 visualisation nodes 2x AMD EPYC 7452, 32-core, 2.35 GHz, 256 GB RAM 1x NVIDIA RTX 6000 GPU	8 visualisation nodes 2x AMD EPYC 7742, 64-core, 2.25 GHz, 2 TB RAM, 8x NVIDIA A40 GPU
Accelerators in total	16x NVIDIA Tesla V100	32x NVIDIA Tesla V100	576x NVIDIA Tesla A100, 2x NVIDIA RTX 6000	11,912x AMD Instinct MI250X, 8x NVIDIA A40
CPU cores in total	48	7,232	106,880	454,784
Storage	30 TB NVMe	29 TB / home, 310 TB / scratch (28 GB/s)	30 TB / home, 1,275 TB / scratch (NVMe, 730 GB/s sequential write performance, 1,198 GB/s sequential read performance)	81 PB / (home + project + scratch) (240 GB/s)
Interconnection	Infiniband EDR 100 Gb/s	Infiniband HDR 200 Gb/s	Infiniband HDR 200 Gb/s	Slingshot-11 200 Gb/s

KARØL1NA

IT4Innovations National Supercomputing Center

VSB – Technical University of Ostrava · Studentska 6231/1B · 708 00 Ostrava · Czech Republic Postal address · 17. listopadu 2172/15 · 708 00 Ostrava · Czech Republic · E-mail: info@it4i.cz

www.it4i.eu

B

A R

B

RA