arnes
povezujemo znanje

EURO
SLING

MREŽA ZNANJA
Ljubljana, 3.–5. december 2024

# HPC Vega – Slovenian EuroHPC supercomputer

## Slovenian Supercomputing Network Day (International)

Dejan Valh
IZUM, Institute of Information Science

# Introduction

IZUM
Institute of Information Science, Maribor, Slovenia

- Public institution, 125 employees (7 for HPC)
- Library automation information system
- Slovenian Current Research Information System
- UNESCO Regional Category II Centre

- Operation & Consortium
- HPC RIVR <-> SLING
- HPC experts from JSI, Arnes and IZUM

EURO SLING

COBISS

COBISS.net

SICRIS

UNESCO
United Nations Educational, Scientific and Cultural Organization

Regional Centre for Library Information Systems and Current Research Information Systems
Regionalni center za knjižnične informacijske sisteme in informacijske sisteme o raziskovalni dejavnosti
IZUM

University of Maribor

FIS Fakulteta za informacijske študije
Faculty of information studies

HPC RIVR

SLING

Slovenian national supercomputing network

arnes

VEGA

Jožef Stefan Institute

# HPC Vega quick facts

- 1st operational EuroHPC JU system
- In production since April 2021
- Performance 6.9 PFLOPS
- Atos Sequana XH2000
- 1020 Compute nodes, Infiniband 100Gb/s
- 18 PB Large Capacity Storage Ceph
- 1 PB High-Performance Storage Lustre
- Power consumption < 1MW
- Power Usage Efficiency, PUE < 1.15
- Hyper-Connected 500 Gb/s

VEGA

TOP500 LIST - JUNE 2021

$R_{max}$ and $R_{peak}$ values are in TFlops. For more details about other fields, check the TOP500 description.

$R_{peak}$ values are calculated using the advertised clock rate of the CPU. For the efficiency of the systems you should take into account the Turbo CPU clock rate where it applies.

| Rank | System | Cores | Rmax (TFlop/s) | Rpeak (TFlop/s) | Power (kW) |
|---|---|---|---|---|---|
| 106 | **VEGA HPC CPU** - BullSequana XH2000, AMD EPYC 7H12 64C 2.6GHz, Mellanox InfiniBand HDR100, Atos IZUM Slovenia | 122,880 | 3,822.0 | 5,367.0 | |
| 134 | **VEGA HPC GPU** - BullSequana XH2000, AMD EPYC 7H12 64C 2.6GHz, NVIDIA A100, Infiniband HDR, Atos IZUM Slovenia | 33,600 | 3,096.0 | 4,680.0 | |

TOP500 The List.

EuroHPC Joint Undertaking

REPUBLIC OF SLOVENIA
MINISTRY OF EDUCATION, SCIENCE AND SPORT

EUROPEAN UNION
COHESION FUND

HPC RIVR

- HPC Infrastructure 17 MIO EUR
- Operation est. costs 30 MIO EUR in 8 years

# Direct Liquid Cooled Data Center

IZUM
Institute of Information Science, Maribor, Slovenia

VEGA
HPC


Data Center RIVR1
Direct Liquid Cooled HPC

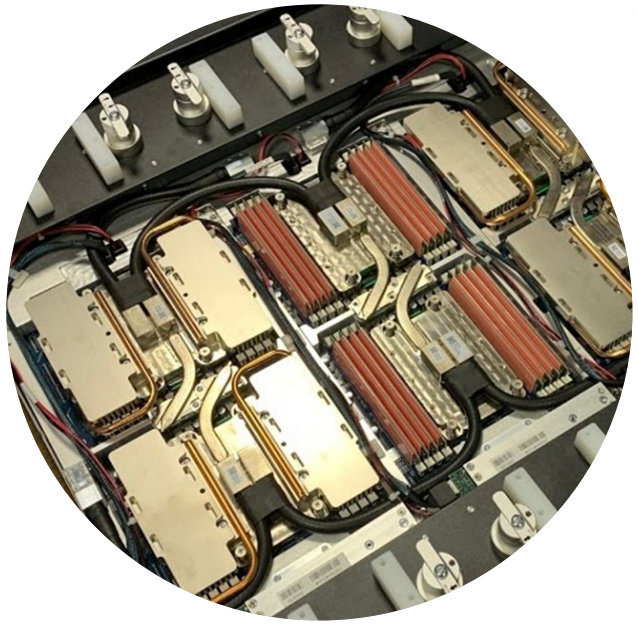

960 CPU nodes
2x AMD 64c, 256 GB / 1 TB

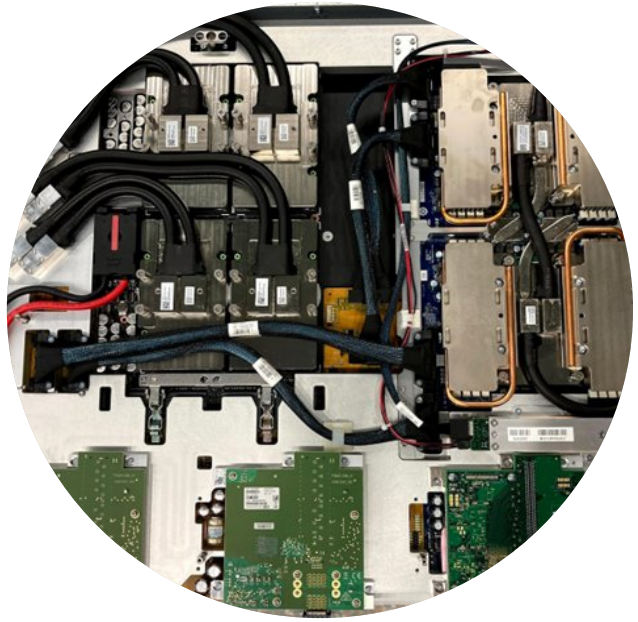

60 GPU nodes
2x AMD 64c, 512 GB
4x NVidia A100, 40 GB

# Share of Resources

**65% of all capacities – national share, divided:**

- up to 70% for Open Access
- up to 20% for Commercial Access (Industry and other)
- from 10% upwards is Hosting Entity (IZUM) reservation
  - Community Access, Urgent Priority Usage
  - Continuous Maintenance

**35% of all capacities – EuroHPC JU share:**

- any EU organization can apply for calls
- independent calls for allocations
- users and projects must be approved by EuroHPC JU

https://doc.vega.izum.si/shares/

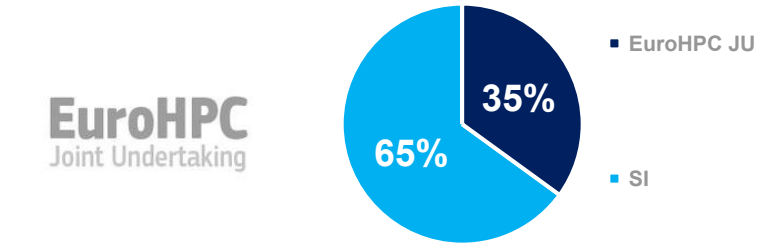


**Resources available for EuroHPC JU**

| Partition | Nodes | Core hours | Storage |
|---|---|---|---|
| CPU standard | 269 | 301.400.064 | 6,3 PB |
| CPU large memory | 67 | 75.125.760 | \ |
| GPU partition | 21 | 23.546.880 | \ |

# Slovenian and EuroHPC JU projects

| SLING calls | Number of proposals | Approved proposals | Ended | Academic and research | Industry |
|---|---|---|---|---|---|
| Temporary projects | 121 | 121 (100 %) | 121 (100 %) | 120 | 1 |
| Test projects | 6 | 6 (100 %) | 3 (50%) | 6 | 0 |
| Development projects | 21 | 21 (100 %) | 0 (0 %) | 20 | 1 |
| Regular research projects | 47 | 47 (100 %) | 0 (0 %) | 46 | 1 |
| Large research projects | 3 | 3 (100 %) + 1 new! | 0 (0 %) | 3 | 0 |
| Commercial projects | 13 | 6 (46 %) | 0 (0 %) | 0 | 6 |
| **Total** | **211** | **204 (97%)** | **124 (61 %)** | **195** | **9** |

https://eurohpc-ju.europa.eu/access-our-supercomputers/eurohpc-access-calls_en

| EHPC call | Number of proposals | Approved proposals | Ended | Academic and research | Industry |
|---|---|---|---|---|---|
| Benchmark projects | 91 | 91 (100 %) | 85 (93 %) | 82 | 9 |
| Development projects | 113 | 113 (100 %) | 60 (53 %) | 99 | 14 |
| Regular research projects | 61 | 43 (70 %) | 31 (72 %) | 39 | 4 |
| **Total** | **265** | **247 (93%)** | **176 (71 %)** | **149** | **27** |

https://www.sling.si/razpisi-za-odprti-dostop-in-rabo-virov/

# Users and projects (accumulated)

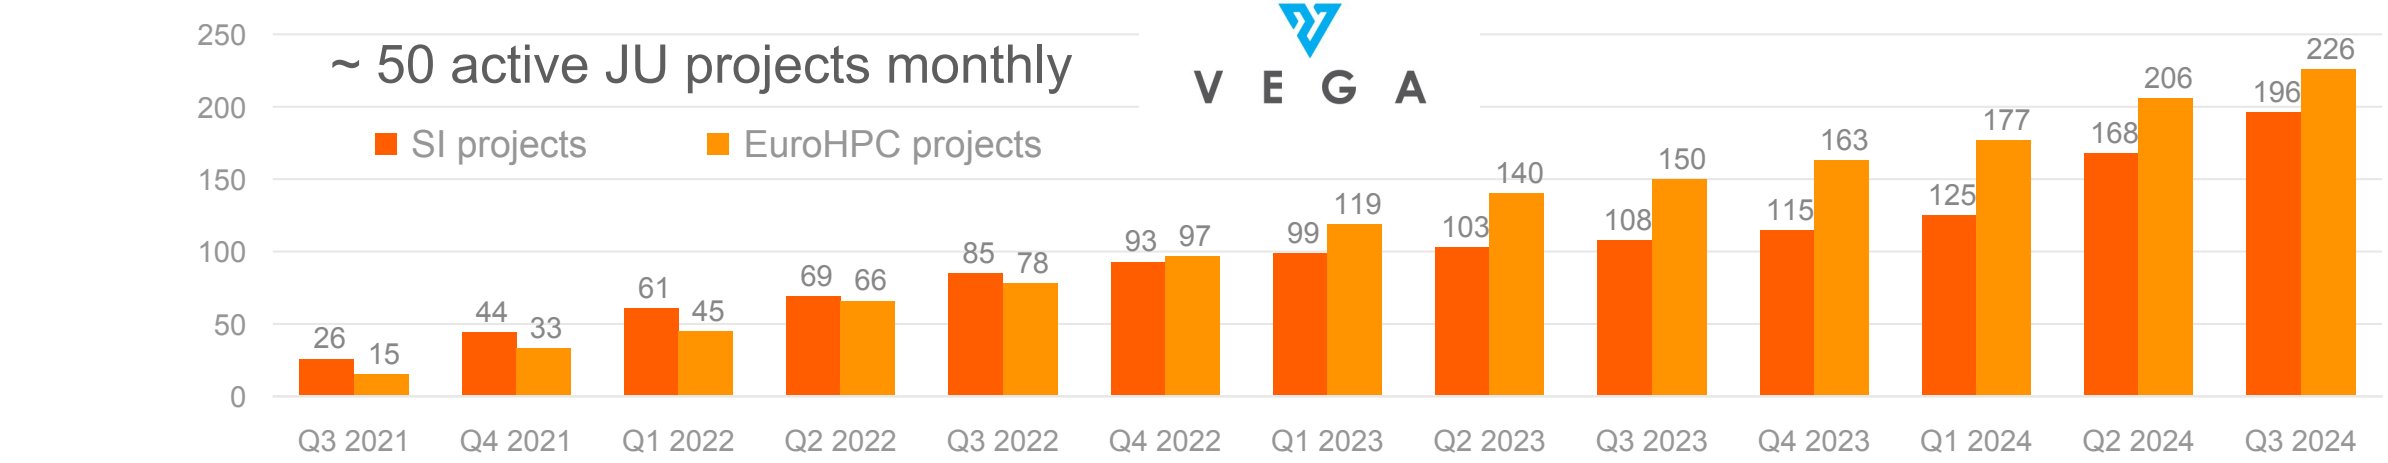


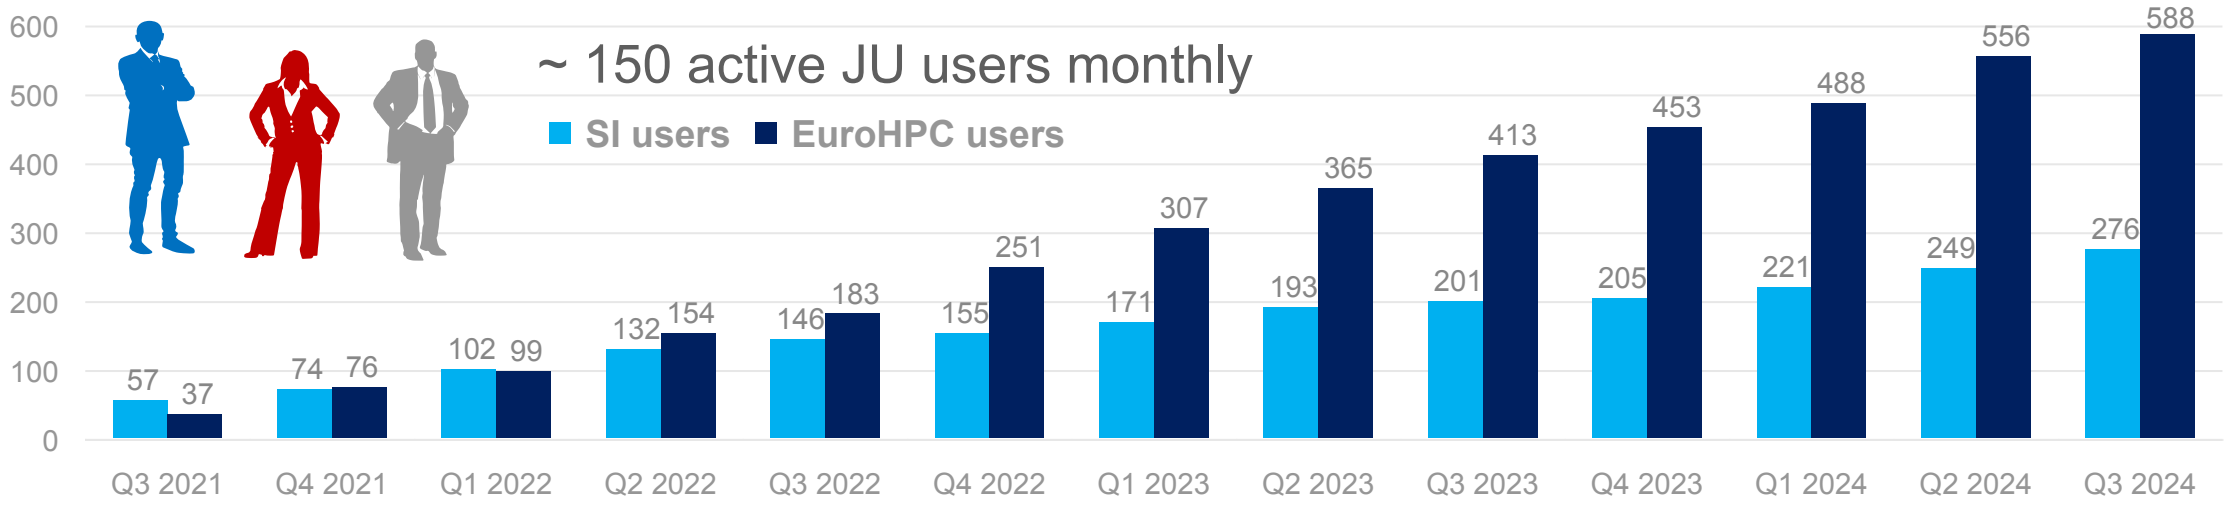

# Regular **EuroHPC JU** projects by the scientific domain

| Domain **Area** | Num. of projects |
|---|---|
| Universe Sciences and Astrophysics | 8 |
| Condensed Matter Physics | 6 |
| Physical and Analytical Chemical Sciences | 5 |
| Products and Processes Engineering | 4 |
| Chemical Sciences and Materials | 3 |
| Engineering | 3 |
| Fundamental Constituents of Matter | 3 |
| Systems and Communication Engineering | 2 |
| Molecular and Structural Biology and Biochemistry | 2 |
| Biochemistry, Bioinformatics and Life sciences | 2 |
| Mathematics and Computer Sciences | 1 |
| Earth System Sciences | 1 |
| Diagnostic Tools, Therapies and Public Health | 1 |
| Genetics, Genomics, Bioinformatics and Systems Biology | 1 |
| Physiology and Medicine | 1 |


6 of 43 projects are involving DL/ML/AI principles

# Large/Regular **domestic** projects by the scientific domain


| Scientific Domain | Num. of projects |
|---|---|
| Physical and analytical chemistry | 12 |
| Computer and Information Science | 11 |
| Condensed matter physics | 7 |
| Materials Engineering | 4 |
| Prevention, diagnosis and treatment of human diseases | 3 |
| Synthetic chemistry and materials | 3 |
| Molecules of Life: Biological Mechanisms, Structures and Functions | 3 |
| Product and Process Engineering | 2 |
| Basics of the substance constituent | 2 |
| Physiology in the field of health, disease and aging | 1 |
| Integrative biology: from genes and genomes to systems | 1 |

19 of 49 projects are involving DL/ML/AI principles

# Beyond Detection: Visual Realism Assessment of Deepfakes

- Artificial intelligence is increasingly applied in all spheres of human society, including science
- Potential dangers associated with AI image generation, deep fakes, need to be addressed
- Scientist are engaged in the detection of artificially manipulated images using an ensemble model with two Convolutional Neural Network (CNN) models (Eva and ConvNext)
- These models (Dragar et al. 2023) have been trained on the DeepFake Game Competition (DFGC) 2022 dataset, where Slovenian scientists secured the third place in the recent DFGC in conjunction with the 2023 International Joint Conference on Biometrics (IJCB 2023)

VEGA

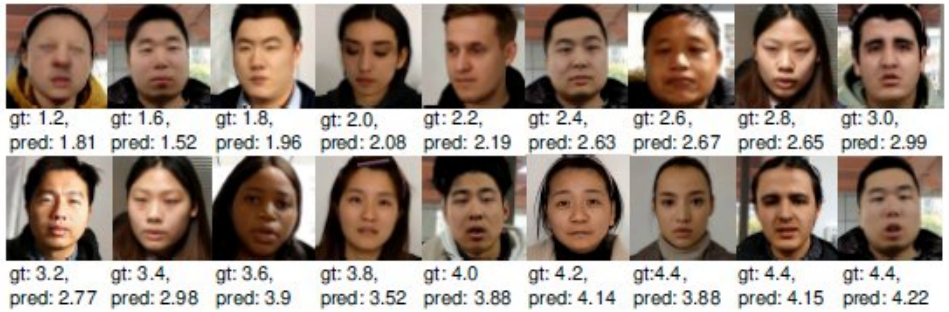


Table 3: Top 3 Models, DFGC-VRA 2023 Competition[1].

| Model | Test Set | PLCC↑ | SRCC↑ | Avg↑ |
|---|---|---|---|---|
| OPDAI | 1 | 0.8578 | 0.8372 | 0.8851 |
| | 2 | 0.9423 | 0.9214 | |
| | 3 | 0.8928 | 0.8592 | |
| HUST | 1 | 0.8117 | 0.7864 | 0.8611 |
| | 2 | 0.9281 | 0.9215 | |
| | 3 | 0.8842 | 0.8348 | |
| Ours | 1 | 0.8091 | 0.7633 | 0.8545 |
| | 2 | 0.9287 | 0.9197 | |
| | 3 | 0.8746 | 0.8318 | |

**Luka Dragar[1], Peter Peer[1], Vitomir Štruc[2], Borut Batagelj[1]**

[1]*Faculty of Computer and Information Science, University of Ljubljana, Večna pot 113, 1000 Ljubljana, Slovenia*
[2]*Faculty of Electrical Engineering, University of Ljubljana, Tržaška cesta 25, 1000 Ljubljana, Slovenia*
*E-mail: luka.dragar3@gmail.com*

# HIDRA2: deep-learning ensemble, sea level and storm tide forecasting

- HIDRA-2 and HIDRA-T application have been developed with the help of HPC Vega (Rus et al. 2023)
- Implementation of the 'deep learning' method, also used in the development of artificial intelligence, was used to develop the HIDRA-2 and HIDRA-T program
- This applications can improve sea level predictions for coastal regions, especially during forecasted storms by 25%, compared to traditional numerical methods
- This applications can independently learn from existing data sources, predict weather extremes and contribute to the protection of people, the environment, infrastructure, and machinery.

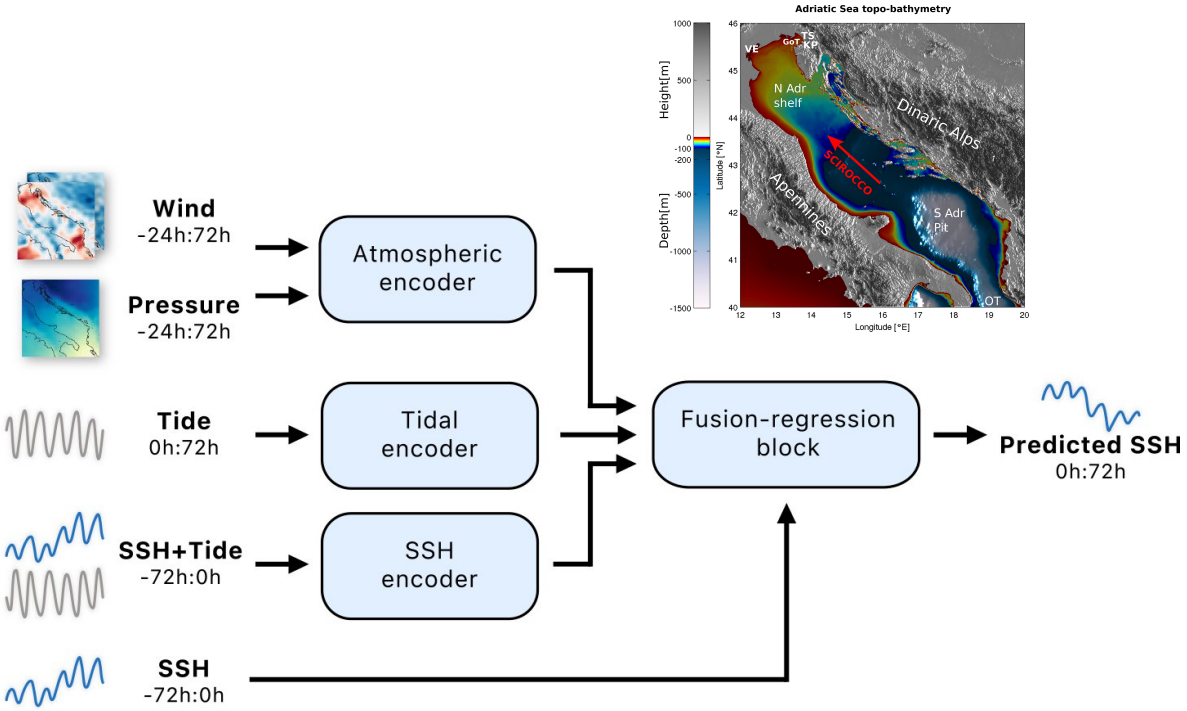



Marko Rus[1], Anja Fettich[2], Matej Kristan[1,★], and Matjaž Ličer[2,3,★]

[1]Faculty of Computer and Information Science, Visual Cognitive Systems Lab, University of Ljubljana, Ljubljana, Slovenia
[2]Slovenian Environment Agency, Office for Meteorology, Hydrology and Oceanography, Ljubljana, Slovenia
[3]National Institute of Biology, Marine Biology Station, Piran, Slovenia
★These authors contributed equally to this work.

**Correspondence:** Marko Rus (marko.rus@fri.uni-lj.si)

# Commercial projects (pay-per-use)

- In silico (Croatia), since the beginning of 2022
- Gorenje, Hisense group (Slovenia) – CFD
- Xlab (Slovenian IT Company) – train AI
- Agenda (Slovenian IT Company) – train AI
- Enerdat-S (Slovenia), Softergee (Slovenian IT Company) - infrastructure
- Algoritmiq (Finland) – Checking QC results
- A few others interested
- E-DIH DIGI-SI (University of Maribor) – vouchers for SMEs, 2.000,00 EUR

In silico
https://insilico.hr/

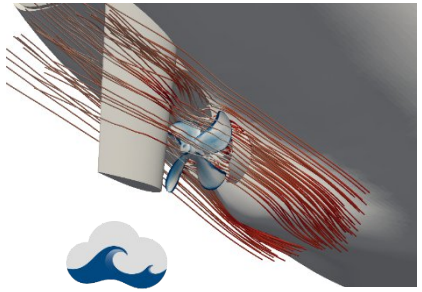

# EuroHPC JU Access Calls

## 2.4 Software and Attributes (1 page)

### *2.4.1 Software*

We use OpenFOAM v2012 as pre-exaFOAM state for the benchmarking; and a later version for the post-exaFOAM benchmark. Alternatives are not possible because the consortium agreed on the versions to be used.

For the workshop simulations, we may use custom improvements of the methods if required.

### *2.4.2 Particular libraries*

OpenFOAM (Finite Volume, SIMPLE method), C++, OpenMPI, cmake, python with numpy and matplotlib, ParaView for remote visualization, user install of SZL server from Tecplot (used for coherent plotting of workshop)

### *2.4.3 Parallel programming*

OpenMPI, heterogenous hardware not used here but targeted by exaFOAM (and benchmarked by other cases).

## 2.6 Performance of Software (Maximum 2 pages)

### *2.6.1 Testing of your code on the requested machine*

The production code was tested on Vega using the Benchmark Access EHPC-BEN-2023B03-016. The test case used here is the WMLES of a 30P30N high-lift configuration. It is as close as possible to the final numerical setup of the workshop case.

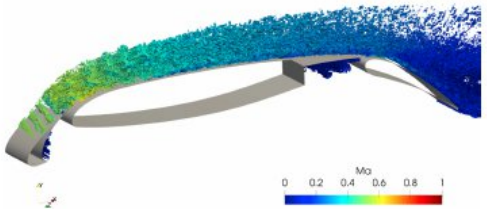


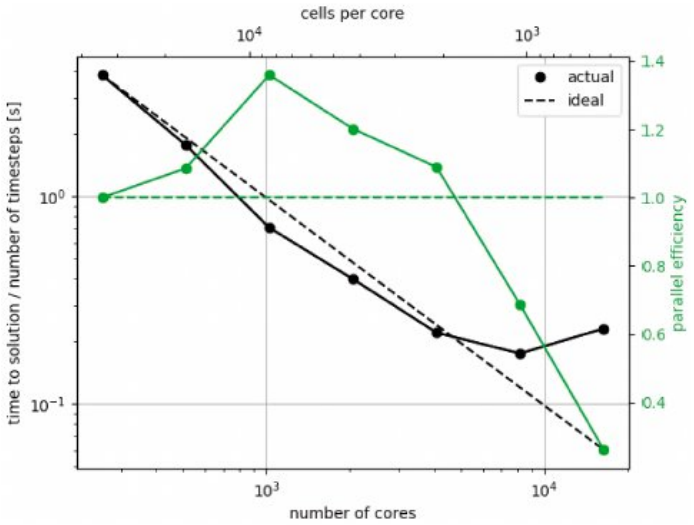


### *2.6.2.3 Time-to-solution*

*Time-to-solution as normalised/averaged per iteration,*

$$T_i^* = \frac{t_i \cdot N_c}{N_e}$$

*AND the normalized total time to solution*

$$T_f^* = \frac{t_f \cdot N_c}{N_e}$$

https://eurohpc-ju.europa.eu/access-our-supercomputers/eurohpc-access-calls_en

### Partitions

Instructions: Not provided

Partitions #1

Partition selection and resources request

Partition name:
Vega CPU

Code(s) used: OpenFOAM

Node-core hrs: Not provided

Requested amount of resources (node hours): 156810

Jobs

Number of jobs simultaneously: 2

Wall clock time of a typical job execution (h): 24

Checkpoints

Are you able to write checkpoint?: Yes

Maximum time between 2 checkpoints (h): 4

Cores/nodes

Minimum # CPU cores: 8192

Average # CPU cores: 30000

Maximum # CPU cores: 64000

# of CPU cores used per node: 128

# of GPUs used per GPU node: Not provided

Memory

Minimum job memory (total usage over all nodes in GB): 4000

Average job memory (total usage over all nodes in GB): 12000

Maximum job memory (total usage over all nodes in GB): 16000

# Availability and usage – CPU

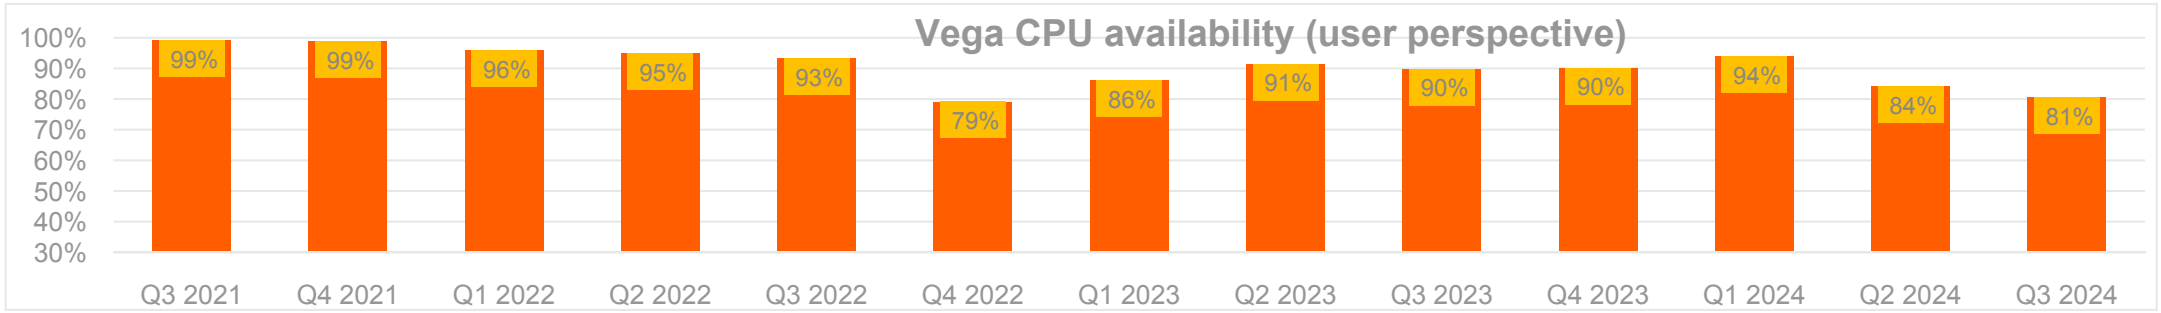


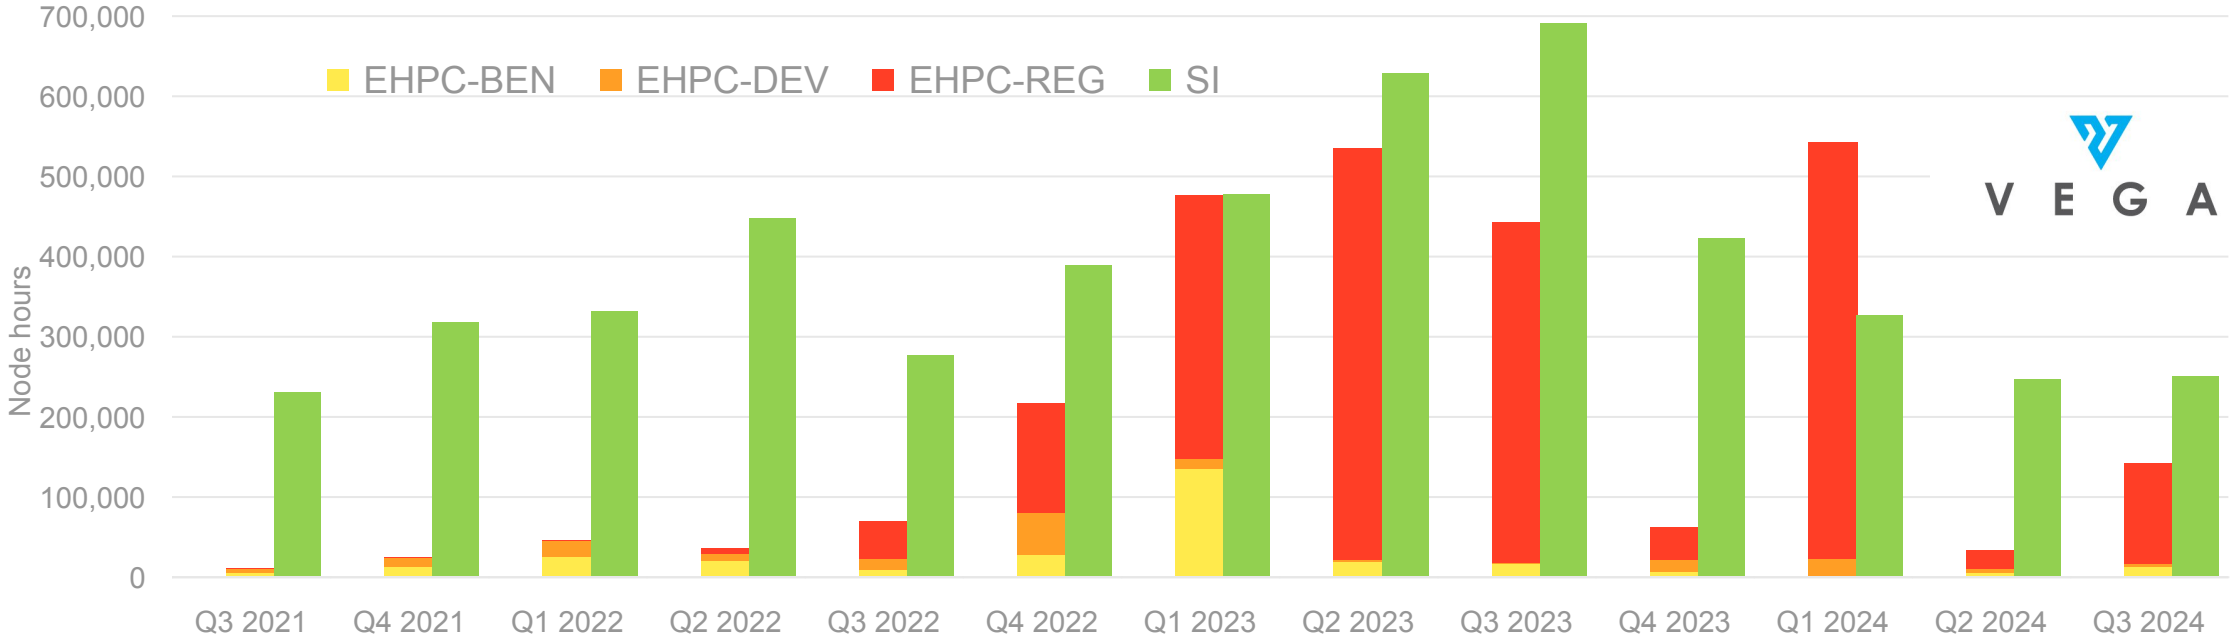

# Availability and usage – GPU

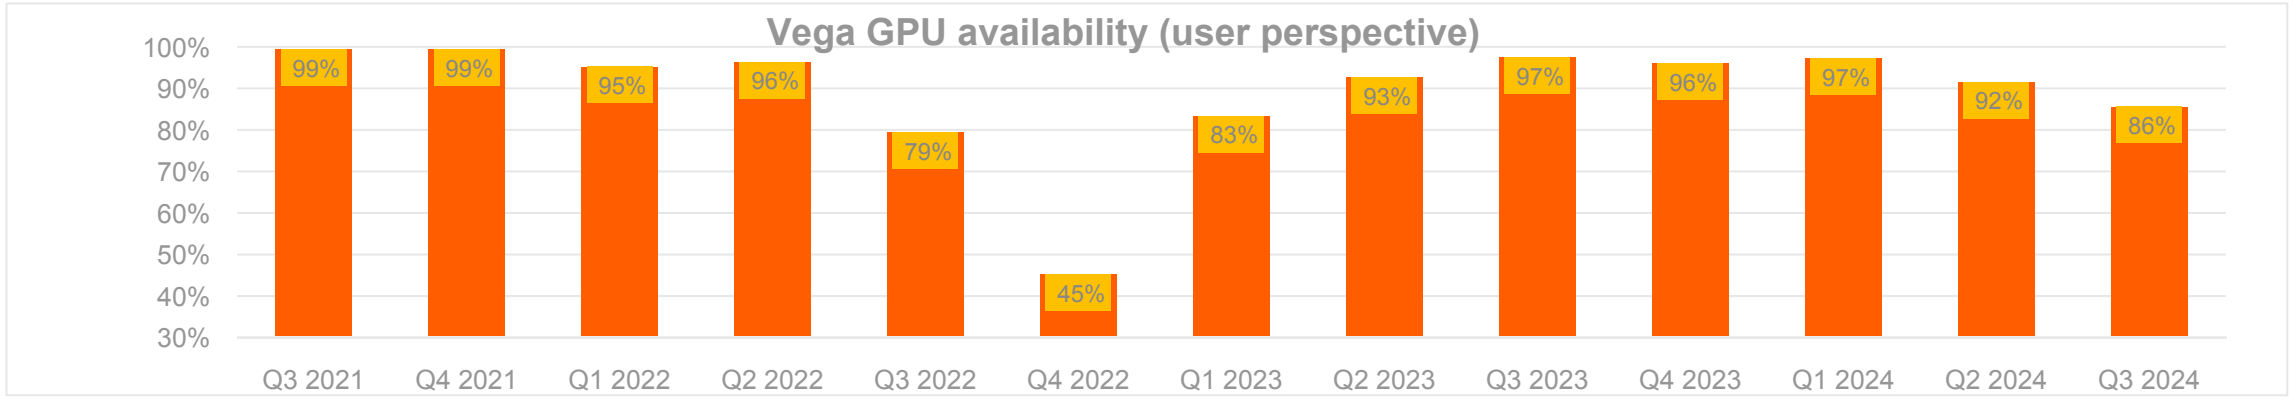


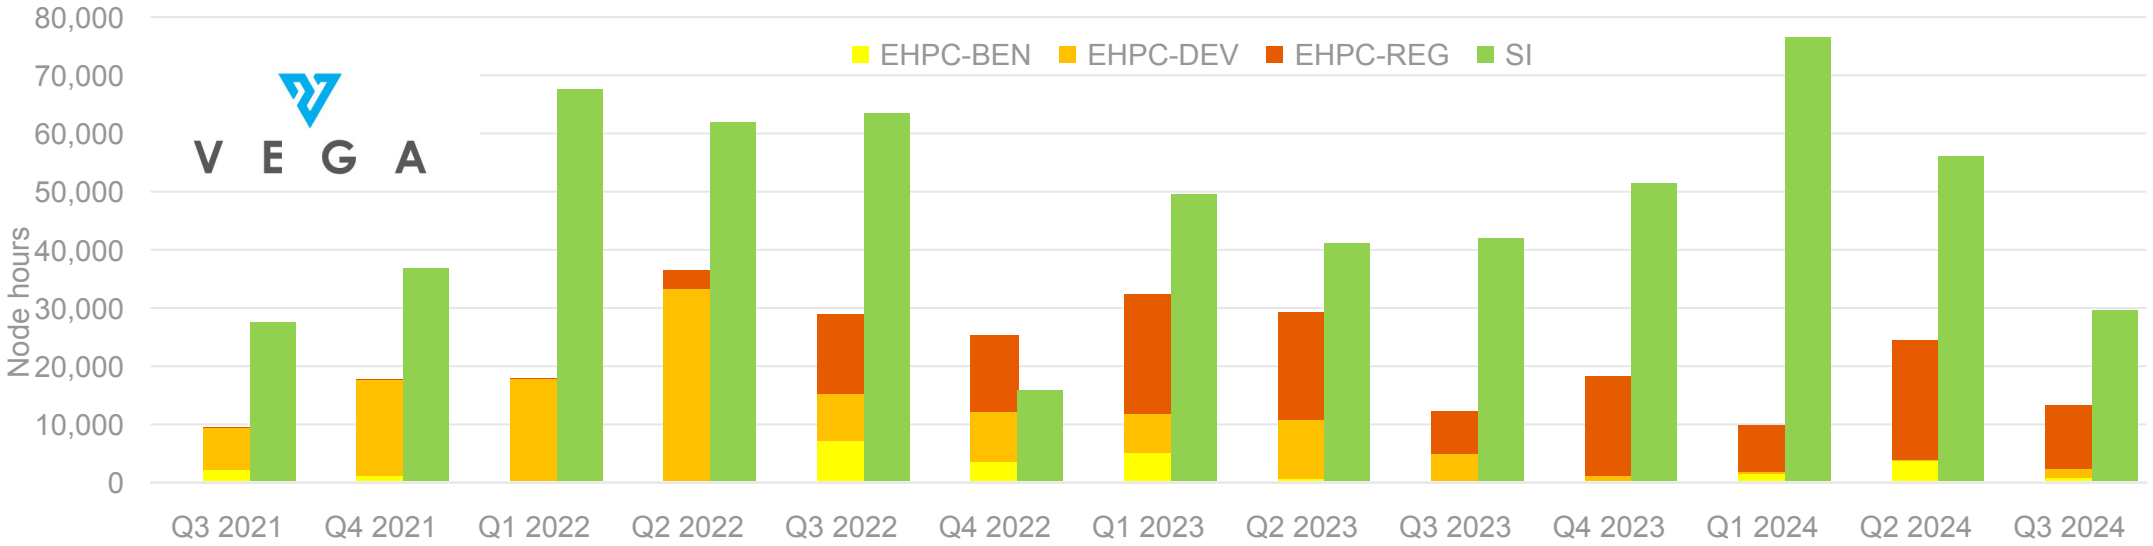

# Number of SLING tickets solved and resolution times

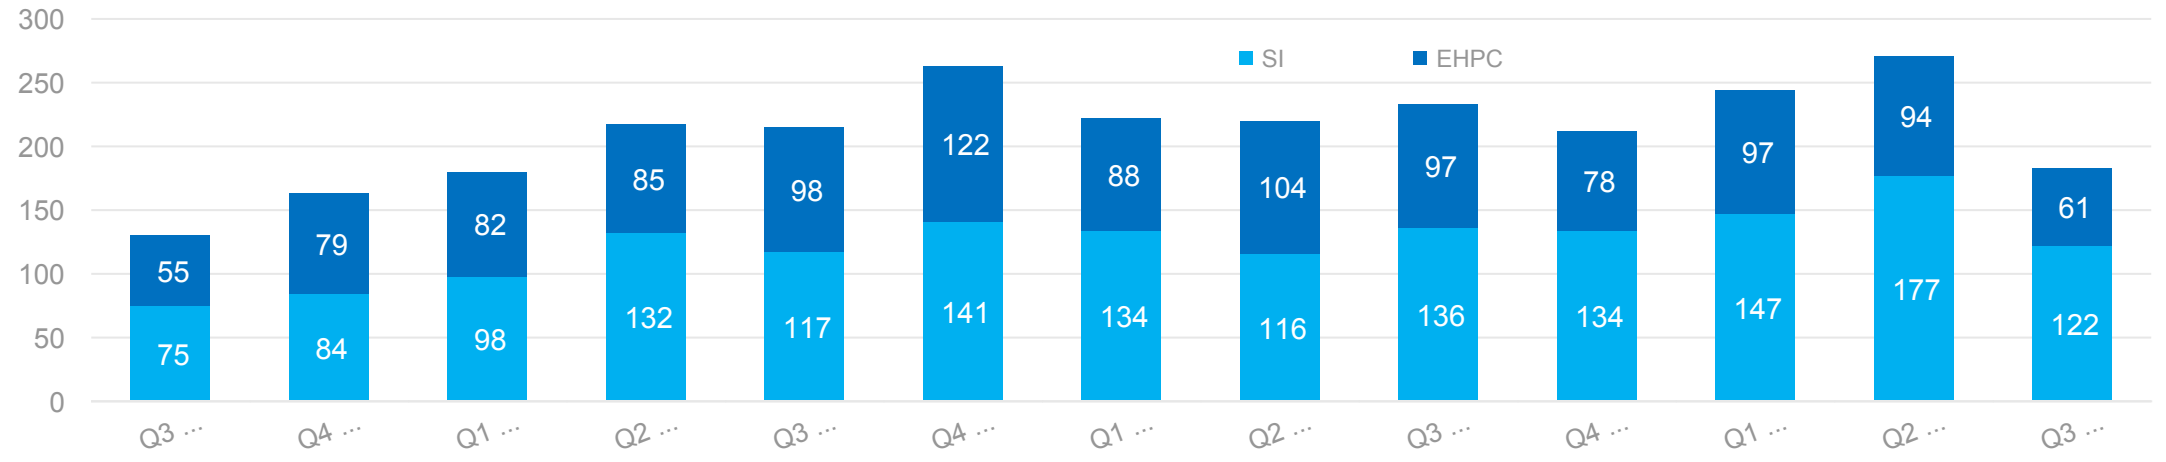


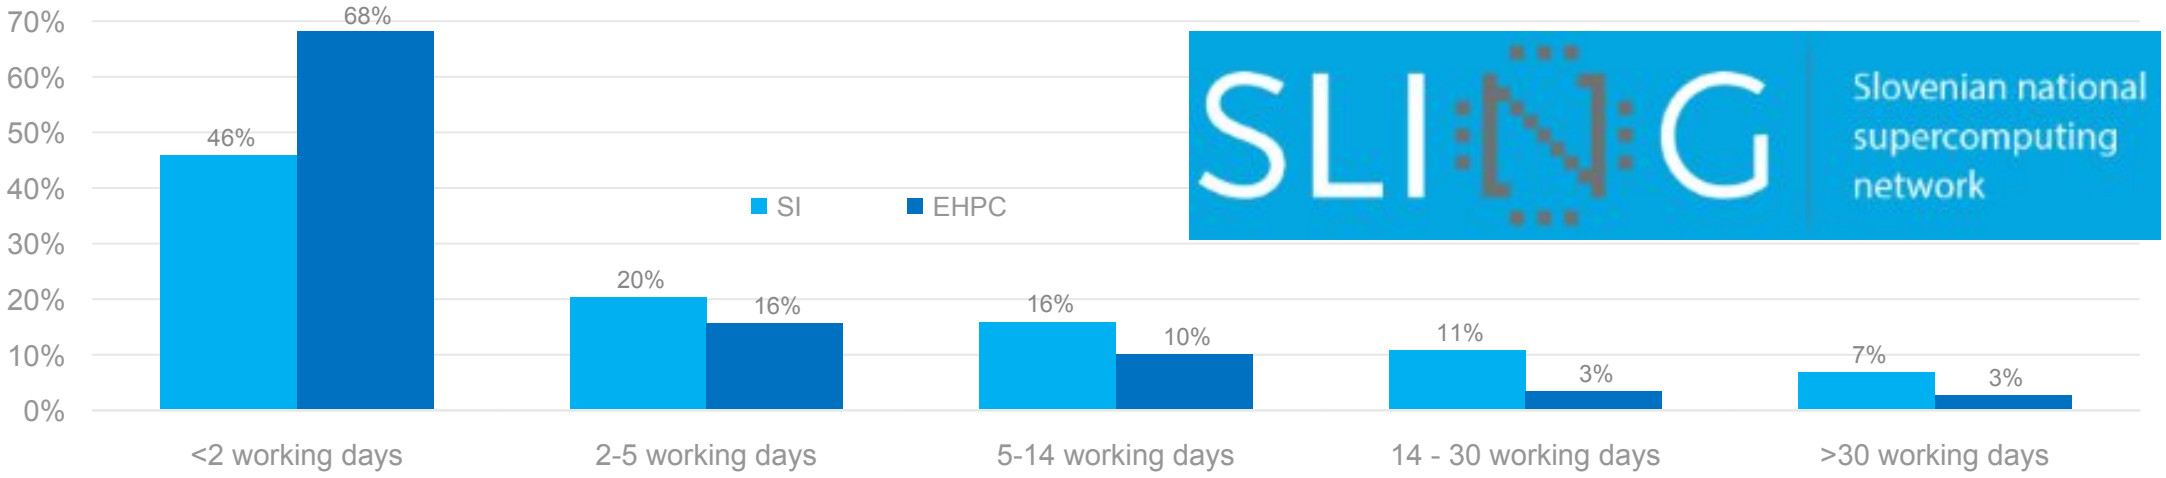


SLING Slovenian national supercomputing network

# EU Projects supported

## Projects (funded)

VEGA

EPICURE
Unlocking European-level HPC Support

SMASH
machine learning for science and humanities postdoctoral programme

interTwin

EUMaster4HPC

## Supporting projects/activities

(not funded)

EURO SLING

MAX DRIVING THE EXASCALE TRANSITION

Multiscale

alt-edic

EESSI
EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

High-level App. Support for Leonardo

CASTIEL2, Container Forum

EVEREST (Experiments for Validation and Enhancement of higher REsolution Simulation Tools, “Safety of operating nuclear power plants and research reactors”)

# Future – New Data center (Power Plant @ Drava river)

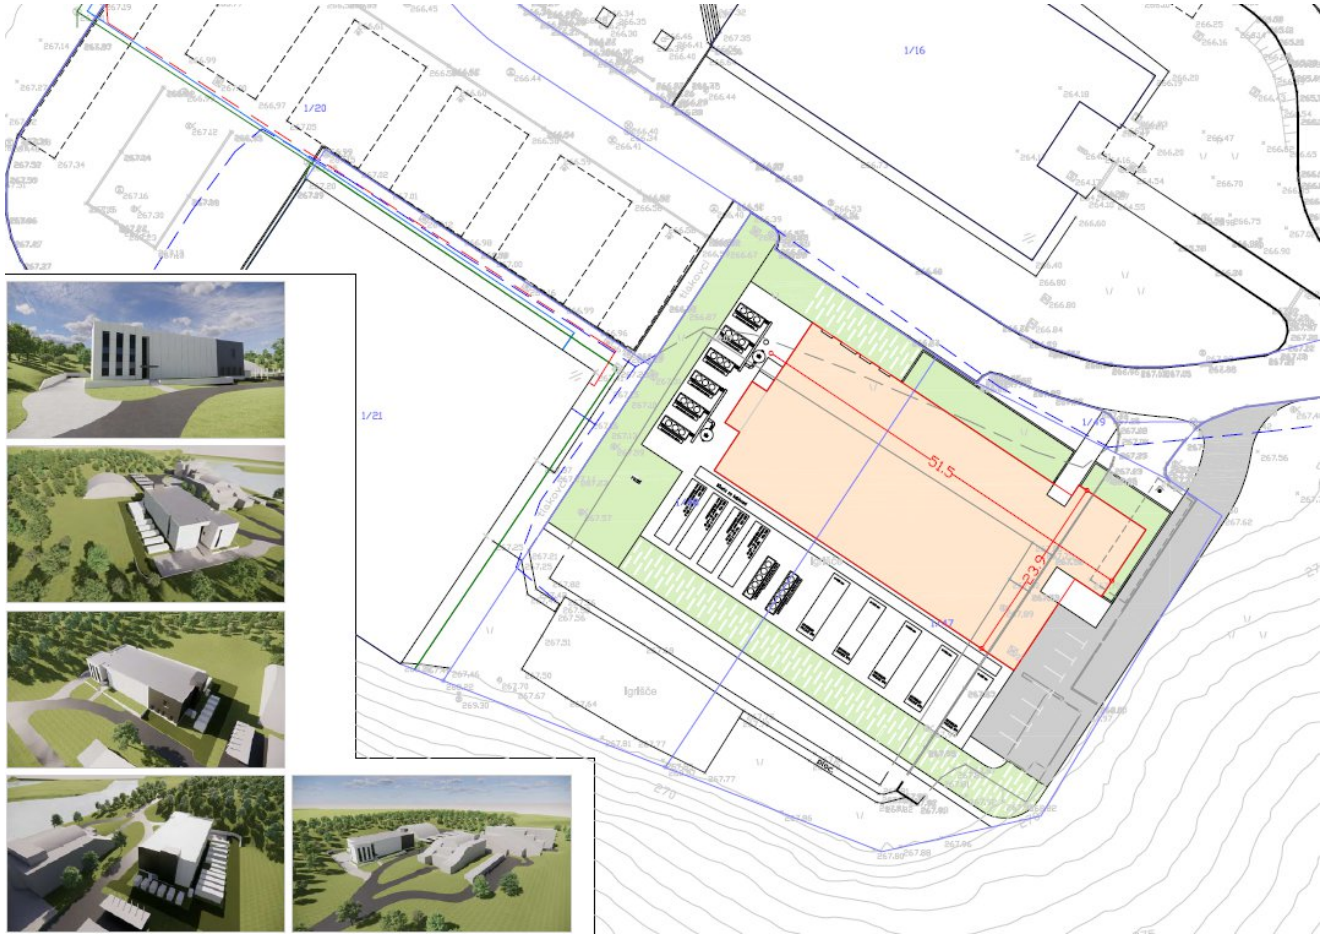


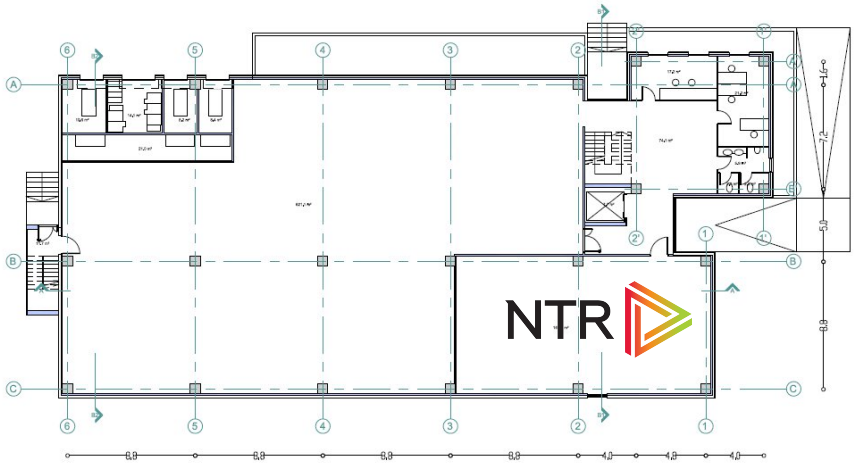


# Special thanks to Dr. Andrej Filipčič, Dr. Jan Jona Javoršek (JSI)

# Funding

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101101903. The JU receives support from the Digital Europe Programme and Germany, Bulgaria, Austria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, Greece, Hungary, Ireland, Italy, Lithuania, Latvia, Poland, Portugal, Romania, Slovenia, Spain, Sweden, France, Netherlands, Belgium, Luxembourg, Slovakia, Norway, Türkiye, Republic of North Macedonia, Iceland, Montenegro, Serbia.
The SLING National Competence Centre is co-funded by the Ministry of Higher Education, Science and Innovation.

## Media sponsor

Računalniške novice
www.racunalniske-novice.com

Co-funded by the European Union

REPUBLIKA SLOVENIJA
MINISTRSTVO ZA VISOKO ŠOLSTVO, ZNANOST IN INOVACIJE