Predstavitev HPC Vega – za podjetja

Superračunalniki za super izzive podjetij

Dejan Valh, Institut informacijskih znanosti, IZUM Maribor

Zmogljivosti

EURO²
SLING

Zmogljivost CPU/GPU

CPU: HPL R_{max} = **3.8 PFLOPS** (R_{peak} 5.4)
GPU: HPL FP64 R_{max} = **3.1 PFLOPS** (R_{peak} 4.7)

CPU GPU

Diskovni sistem za trajno hrambo

kapaciteta 18 PB
prepustnost >200 GB/s
datotečni sistem Ceph

ceph 18 PB

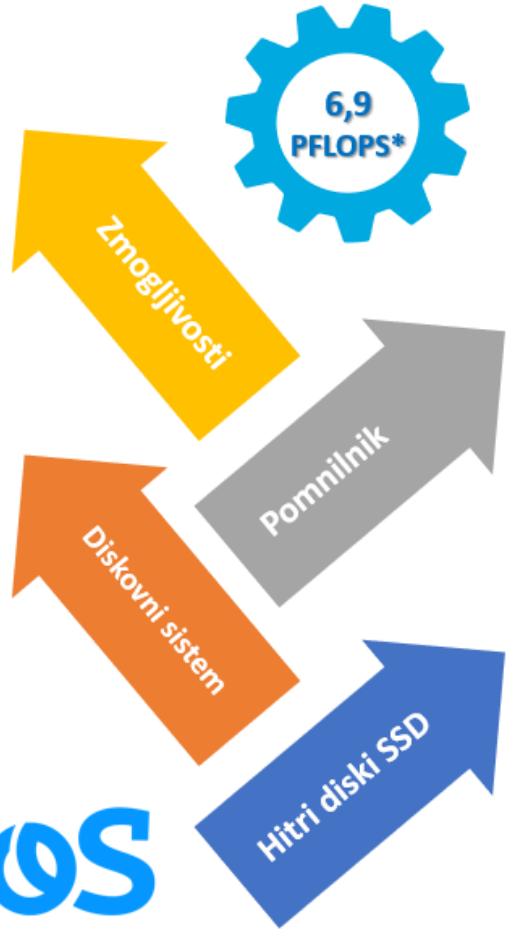

*FLOPS = operacij s plavajočo vejico

Pomnilnik

Vozlišče CPU: 256 GB ali 1 TB,
(2 ali 8 GB na processor)
Vozlišče GPU: 512 GB,
4 GB na procesor
40 GB HBM2 na grafični pospeševalnik

Visoko zmogljiv diskovni sistem

kapaciteta 1 PB
prepustnost >400 GB/s
datotečni sistem Lustre

lustre 1 PB

Omrežje IB/Eth 100 Gb/s

VEGA

Povzetek: 100 m², 1.1 MW, 123k jeder, 2-8 GB/jedro, 240 GPU-jev, 6.9 PFLOPS, 19 PB, 500 Gbps

Particija CPU

1 256GiB/1024GiB System Memory
16 16GiB DIMM DDR4 Synchronous Registered (Buffered) 3200 MHz (0.3 ns)
2 AMD EPYC 7H12 64-Core Processor
1 MT28908 Family [ConnectX-6]
1 Micron_7300_MTFDHBG1T9TDF
1 1920GB NVMe disk
2 I350 Gigabit Backplane Connection

960 BullSequana XH2410 vozlišč

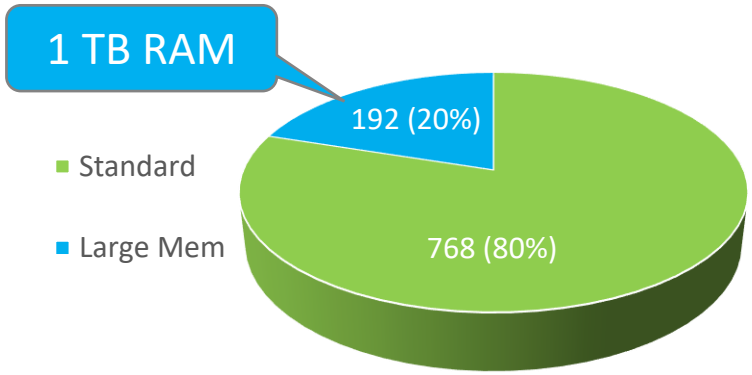

CPU vozlišča

191 CPU in 48 vozlišč LM za potencialno uporabo podjetjem!

```
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                256
On-line CPU(s) list:   0-255
Thread(s) per core:    2
Core(s) per socket:    64
Socket(s):             2
NUMA node(s):          8
Vendor ID:             AuthenticAMD
CPU family:            23
Model:                 49
Model name:            AMD EPYC 7H12 64-Core Processor
Stepping:              0
CPU MHz:               2927.281
CPU max MHz:           3313.4761
CPU min MHz:           1500.0000
BogoMIPS:              5200.37
Virtualization:        AMD-V
L1d cache:             32K
L1i cache:             32K
L2 cache:              512K
L3 cache:              16384K
```

280W

Particija GPU

1 **512GiB** System Memory
16 32GiB DIMM DDR4 Synchronous Registered (Buffered) 3200 MHz (0.3 ns)
2 AMD EPYC 7H12 64-Core Processor
4 MT28908 Family [ConnectX-6]
1 Micron_7300_MTFDHBG1T9TDF
1 1920GB NVMe disk
2 I350 Gigabit Backplane Connection

60 BullSequana XH2415 vozlišč

15 GPU vozlišč za potencialno rabo podjetij

VEGA

Nvidia Datacenter GPU	Nvidia A100
GPU codename	GA100
GPU architecture	Ampere
Launch date	May 2020
GPU process	TSMC 7nm
Die size	826mm2
Transistor Count	54 billion
FP64 CUDA cores	3,456
FP32 CUDA cores	6,912
Tensor Cores	432
Streaming Multiprocessors	108
Peak FP64	9.7 teraflops
Peak FP64 Tensor Core	19.5 teraflops
Peak FP32	19.5 teraflops
Peak FP32 Tensor Core	156 teraflops/312 teraflops*
Peak BFLOAT16 Tensor Core	312 teraflops/624 teraflops*
Peak FP16 Tensor Core	312 teraflops/624 teraflops*
Peak INT8 Tensor Core	624 teraflops/1,248 TOPS*
Peak INT4 Tensor Core	1,248 TOPS/2,496 TOPS*
Mixed-precision Tensor Core	312 teraflops/624 teraflops*
Max TDP	400 watts

*Effective TOPS / TFLOPS using the new Sparsity feature

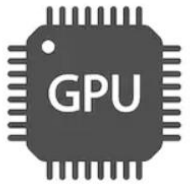

https://askgeek.io/en/gpus/vs/NVIDIA_A100-SXM4-40-GB-vs-NVIDIA_GeForce-RTX-3070

Štirje pospeševalniki na vozlišče GPU

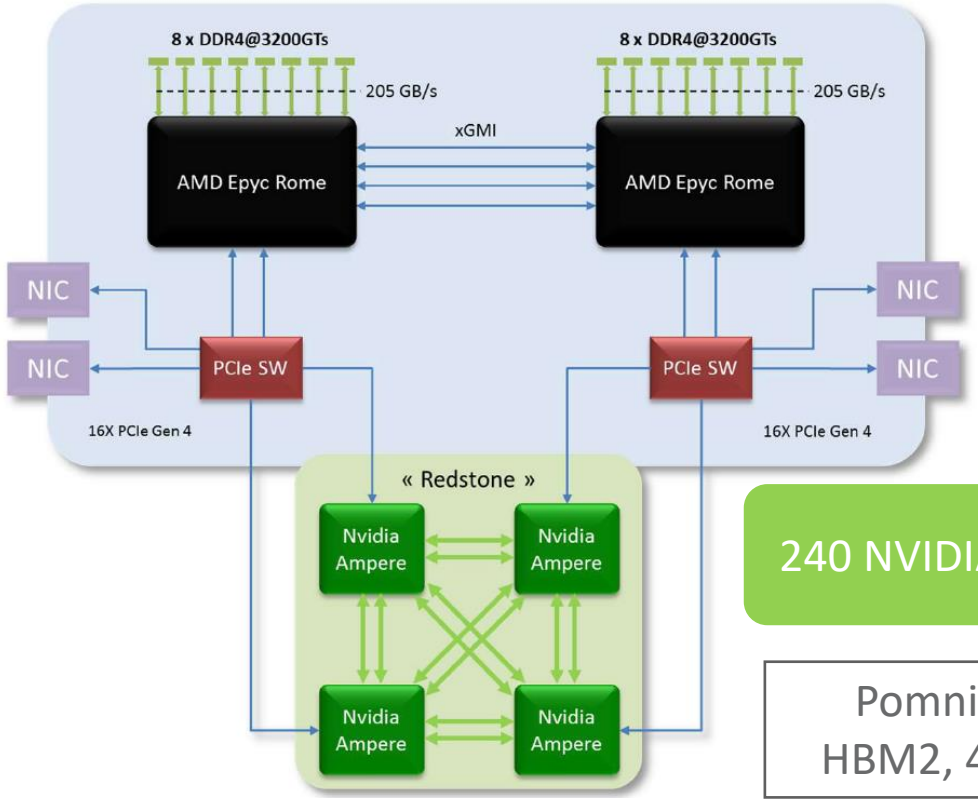

240 NVIDIA A100

Pomnilnik HBM2, 40 GB

60 pospeševalnikov GPU za potencialno rabo podjetij

```
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 535.54.03              Driver Version: 535.54.03    CUDA Version: 12.2|
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|                               |                      |               MIG M. |
|===============================+======================+======================|
|   0  A100-SXM4-40GB      On   | 00000000:03:00.0 Off |                    0 |
| N/A   46C    P0    60W / 400W |      3MiB / 40536MiB |      0%      Default |
|                               |                      |             Disabled |
+-------------------------------+----------------------+----------------------+
|   1  A100-SXM4-40GB      On   | 00000000:44:00.0 Off |                    0 |
| N/A   59C    P0   224W / 400W |    532MiB / 40536MiB |     95%      Default |
|                               |                      |             Disabled |
+-------------------------------+----------------------+----------------------+
|   2  A100-SXM4-40GB      On   | 00000000:84:00.0 Off |                    0 |
| N/A   51C    P0    87W / 400W |    730MiB / 40536MiB |      0%      Default |
|                               |                      |             Disabled |
+-------------------------------+----------------------+----------------------+
|   3  A100-SXM4-40GB      On   | 00000000:C4:00.0 Off |                    0 |
| N/A   43C    P0    58W / 400W |      0MiB / 40536MiB |      0%      Default |
|                               |                      |             Disabled |
+-------------------------------+----------------------+----------------------+
```

Infrastruktura je primerna za učenje modelov, ne za interaktivne seje.

Vozlišča za prijavo in virtualizacijo

8 vozlišč za prijavo (login node)

- 256 GB DDR4
- enaki procesorji kot v CPU vozliščih
- enaka grafična kartica kot v GPU vozliščih
- dva lokalna diska 7.6 TB NVMe
- 1 x 100 Gb/s, Infiniband HDR100
- 2 x 100 Gb/s, Ethernet

30 vozlišč za virtualne gostitelje

- 512 GB pomnilnika DDR4 na vozlišče
- dva procesorja AMD Epyc 7502
  (32 jeder, 2.5 GHz, 180 W)
- dva lokalna diska 7.6 TB NVMe
- 1 x 100 Gb/s, Infiniband HDR100
- 2 x 100 Gb/s, Ethernet

2 vozlišči za prijavo za potencialno rabo podjetij + namenski virtualni strežniki

Dva diskovna sistema

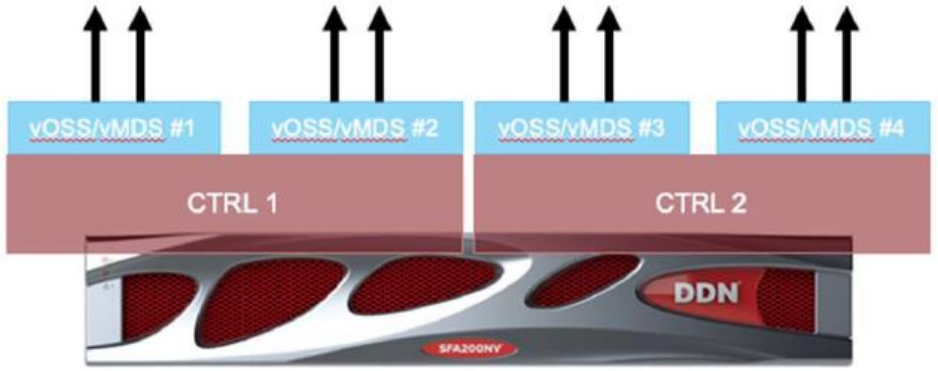

VEGA

10 sistemov DDN – visokozmogljivi sistem HPST

- 23 diskov 6.4 TB NVMe na sistem
- uporabne kapacitete 111 TB na sistem
- 4 vgrajeni virtualni strežniki za Lustre
- povezljivost 8 x 100 Gb/s, Infiniband HDR100

Prepustnost obdelave podatkov je 400 GB/s.

61 vozlišč za shranjevanje podatkov LCST

- vsa vozlišča z 256 GB pomnilnika DDR4
- dva procesorja Intel 4214R (12 jeder, 2.4 GHz, 100 W)
- dva diska 6,4 TB NVMe za sistem
- 24 diskov 16 TB
- povezljivost 2 x 25 Gb/s Ethernet
  (do privatnega internega omrežja,
  naprej večkrat 100 Gb/s)

Prepustnost obdelave podatkov je 200 GB/s.

4,5 PB LCST in 250 TB HPST za potencialno uporabo podjetjem

Deleži namenjeni za rabo podjetjem

VEGA

Vir	Viri: Vsi	Viri: SI delež (65%)	Viri: EuroHPC delež (35%)	Odprti dostop EuroHPC: Regularni poziv (industry track 20% od 70%)	Komerciala: SI (20%)	Komerciala: EuroHPC (20%)	SKUPAJ
CPU [vozlišča]	768	499	269	38	100	54	191
LM [vozlišča]	192	124	68	10	25	14	48
GPU [vozlišča]	60	39	21	3	8	4	15
Sistem LCST [PB]	18	11,7	6,3	0,882	2,34	1,26	4,48
Sistem HPST [TB]	1000	650	350	49	130	70	249

Hitri odprti dostop za podjetja preko EuroHPC JU:

	Vega CPU [node hours]	Vega GPU [node hours]	Link
BENCHMARK	5.000	400	https://eurohpc-ju.europa.eu/eurohpc-ju-call-proposals-benchmark-access_en
DEVELOPMENT	10.000	1.000	https://eurohpc-ju.europa.eu/eurohpc-ju-call-proposals-development-access_en

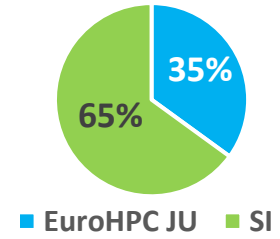

Projekti na HPC Vega

- V odprtem dostopu EuroHPC JU (SME):

Benchmark	Development	Regular
7	8	2

VEGA

In silico
https://insilico.hr/

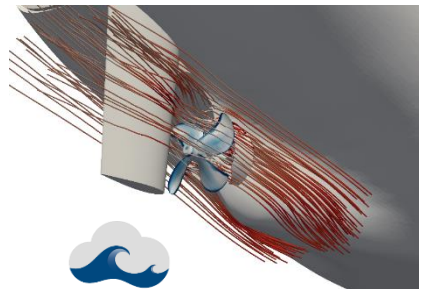

- Komercialni dostop:
  - In silico (Hrvaška), od leta 2022
  - Gorenje – Hisense group – CFD
  - Xlab – učenje jezikovnih modelov
  - Agenda – učenje modelov umetne inteligence

gorenje a Hisense company

XLAB

AGENDA OPEN SYSTEMS

VABLJENI NA OGLED V ŽIVO!

https://www.izum.si/virtualni-sprehod/VEGA.html

Hvala!

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

Načini dostopov do virov za podjetja

Superračunalniki za super izzive podjetij

Dejan Valh, Institut informacijskih znanosti, IZUM Maribor

20. oktober 2023

Kaj pomeni odprti dostop do virov

- Je ***BREZPLAČEN*** (fundiran vnaprej), vendar so zahteve jasne:
  1. **Priznanje** (zahvala) financerjem (npr. omemba HPC Vega, projekta HPC RIVR, zahvala EuroHPC JU) v vseh predstavitvah, objavah, promocijah itd.
  2. **Priprava poročila in objava rezultatov raziskave** (npr. objava modelov, rešitev, postopkov, procesov)
  3. Sodelovanje v **diseminacijskih aktivnostih in promociji** infrastrukture in financerjev
- Kako podjetja pridobijo vire v odprtem dostopu:
  - Preko pozivov EuroHPC JU za deleže EuroHPC na **vseh sistemih EuroHPC**
  - Na nacionalnem deležu HPC Vega in ostalih preko **sodelovalnih projektov**:
    - **raziskovalne narave** npr. projekt, ki ga izvaja fakulteta/institut v sodelovanju s podjetjem ali pa da je npr. zaposleni v podjetju vpisan na doktorski študij ali pa da sodeluje v centrih odličnosti EuroHPC (CoE) ali pa da sodeluje na projektih EuroHPC JU (npr. interTwin).
    - **za izobraževalne in testne (uvajalne) namene**, npr. preko EuroCC SLING, E-DIH, DIGI-SI

Prijava na pozive EuroHPC JU

Project Application

The Project

Project name: Skin permeability

Research field title: LS1 Molecular and Structural Biology and Biochemistry

Research field sub-title: LS1_1 Molecular biology and interactions

Research field share (%): 100

Additional keywords:

Skin permeability prediction
Topical formulation development

Project summary (abstract):

A molecular level understanding of skin permeation may rationalize and streamline product development and improve quality and control, of transdermal and topical drug delivery systems. We have developed an atomistic model of the main barrier in human skin. We use that model to predict drug permeability using molecular dynamics (MD) simulations (with GROMACS). We have also shown that we can predict the effect of permeation enhancers in pharmaceutical formulations. This will help us improve topical drug delivery.

Permeability calculations, using MD simulations, through barrier structures with long correlation times, such as the lipid barrier in skin, require much sampling. The compute capabilities of HPC significantly enhance the throughput, making it possible to evaluate the properties of complex pharmaceutical transdermal, or dermal, formulations consisting of multiple components.

Proposal for civilian purposes: Proposal for civilian purposes

Resubmission type: no

ID of the original project for resubmission (if applicable): Not provided

Track selection: Industry Access Track

Partitions

Partition name: Vega GPU

Code used: GROMACS
Requested amount of resources 1st year in core hours: 3900000
Number of jobs simultaneously: 2
Wall clock time of a typical job execution (h): 48
Are you able to write checkpoint?: yes
Maximum time between 2 checkpoints (h): 1
Minimum # cores: 384
Average # cores: 384
Maximum # cores: 384
Minimum job memory (total usage over all cores GB): 20
Average job memory (total usage over all cores GB): 20
Maximum job memory (total usage over all cores GB): 20
Maximum amount of SCRATCH needed at a time (GB): 20
Maximum amount of WORK needed at a time (GB): 100

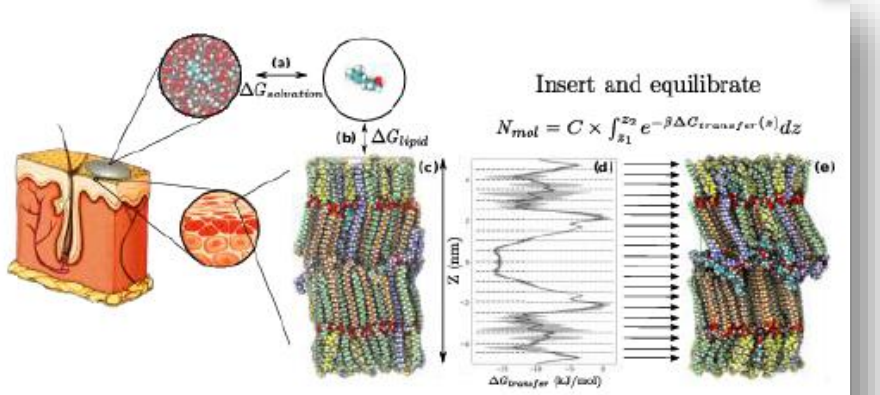

Figure 1

2.5 Description of the results obtained from the scientific point of view, future perspectives, benefits to the society and the benefits of using computer resources

Please fill in the information in the box below (maximum 1000 words).

Understanding how chemical permeation enhancers (PEs) impact skin to improve transdermal delivery of drugs has great implications for clinical medicine, as replacing oral, or intravenous, with topical, drug administration, would give physicians better control over medication and dosage levels. Due to the inability of current in vitro methods to predict the effects of PEs on transdermal drug transport, complementary in silico-methods are requested.

Using molecular dynamics (MD) simulation we have predicted the effect of PEs on drug permeation through skin. Qualitative agreement between calculated- and in vitro determined permeability coefficients is good, and MD simulations are able to reproduce PE permeability enhancement ratio (ER) rankings.

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

O komercialnem (tržnem) dostopu

- Rabo virov je treba plačati (pay-per-use)
- Splošna **omejitev** za tržno dejavnost na raziskovalni infrastrukturi je **20%** (regulacija EU)
- Brez omejitev so komercialni ponudniki
- Oblike in načini izvedbe so različni glede na potrebe podjetij in njihovo **zrelost za uporabo superračunalnikov**
- Kdo ponuja storitve v Sloveniji:
  - Podjetje *Arctur* kot komercialni ponudnik za **razne vrste projektov (razvojne)**
  - *Univerza v Mariboru* na raziskovalni infrastrukturi HPC Maister za **projekte z manj viri**
  - *IZUM* na raziskovalni infrastrukturi HPC Vega za projekte, kjer se potrebuje **veliko računskih virov**

Komercialna raba HPC Vega (1)

- Za podjetja, ki imajo **večje potrebe po računskih virih** (npr. ko jim ostala infrastruktura ne zadostuje in imajo dovolj znanja in izkušenj)
- Na voljo je **20% od nacionalnega deleža** (13% HPC Vega) – če bi to presegli, potem je na voljo **tudi 20% od deleža EuroHPC** (7%), cena je za oba dela enaka (https://si-vegadocs.vega.izum.si/cenik/)!
- Primerna so podjetja, ki imajo **več izkušenj pri rabi HPC**:
  - Znanje in izkušnje morajo pridobiti preko izobraževanj EuroCC SLING in ostalih
  - Razvojno pomoč, svetovanje in testiranje za dvig zrelosti lahko pridobijo pri komercialnih ponudnikih, kot je Arctur
  - Podpora HPC Vega je v prvi vrsti namenjena izkušenim raziskovalnim skupinam

Komercialna raba HPC Vega (2)

- Postopek do uspešne komercialne rabe na HPC Vega:
  1. Inicialni sestanek – ugotovi se primernost za računanje na HPC Vega
  2. Registracija: https://www.sling.si/forms/komercialni/
  3. Uvajalno obdobje (sodelovanje s podporo SLING):
     - Namestitev programske opreme, prilagajanje okolja (npr. dodatni virtualni strežniki)
     - Kopiranje podatkov in zagon testnih poslov na razvojni particiji
     - Če je uvajanje uspešno, sledijo nadaljnji koraki
  4. Podpis sporazuma oz. pogodbe, če je to zahteva s strani podjetja (na voljo predloga)
  5. Podpis sporazuma o ne razkritju podatkov, če je to zahteva podjetja (na voljo predloga)
  6. Plačilo letnega pavšala (letna obnovitev)
  7. Začetek produkcijske rabe sistema
  8. Izdaja mesečni računov glede na porabo za en mesec nazaj

Hvala!

POGLEJTE HPC Industry Summit:
https://www.youtube.com/watch?v=1VdejU1PIpk
https://www.youtube.com/watch?v=CgQET4qIoCY

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