EuroHPC USER DAY

22 – 23 October 2024

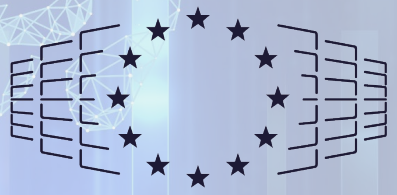

EuroHPC
Joint Undertaking

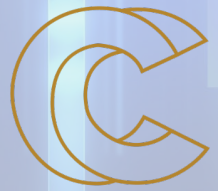

EURO
NETHERLANDS

Usage of the EuroHPC JU supercomputers

STATE OF PLAY ON AVAILABLE SYSTEMS & RESULTS OF EUROHPC JU ACCESS CALLS

EuroHPC JU User Day 2024

Klara Meštrović & Dora Marton

Overview

EuroHPC Joint Undertaking

INFRASTRUCTURE PROCUREMENT

SUPERCOMPUTERS

ACCESS TO SUPERCOMPUTERS

ACCESS MODES

PEER-REVIEW PROCESS

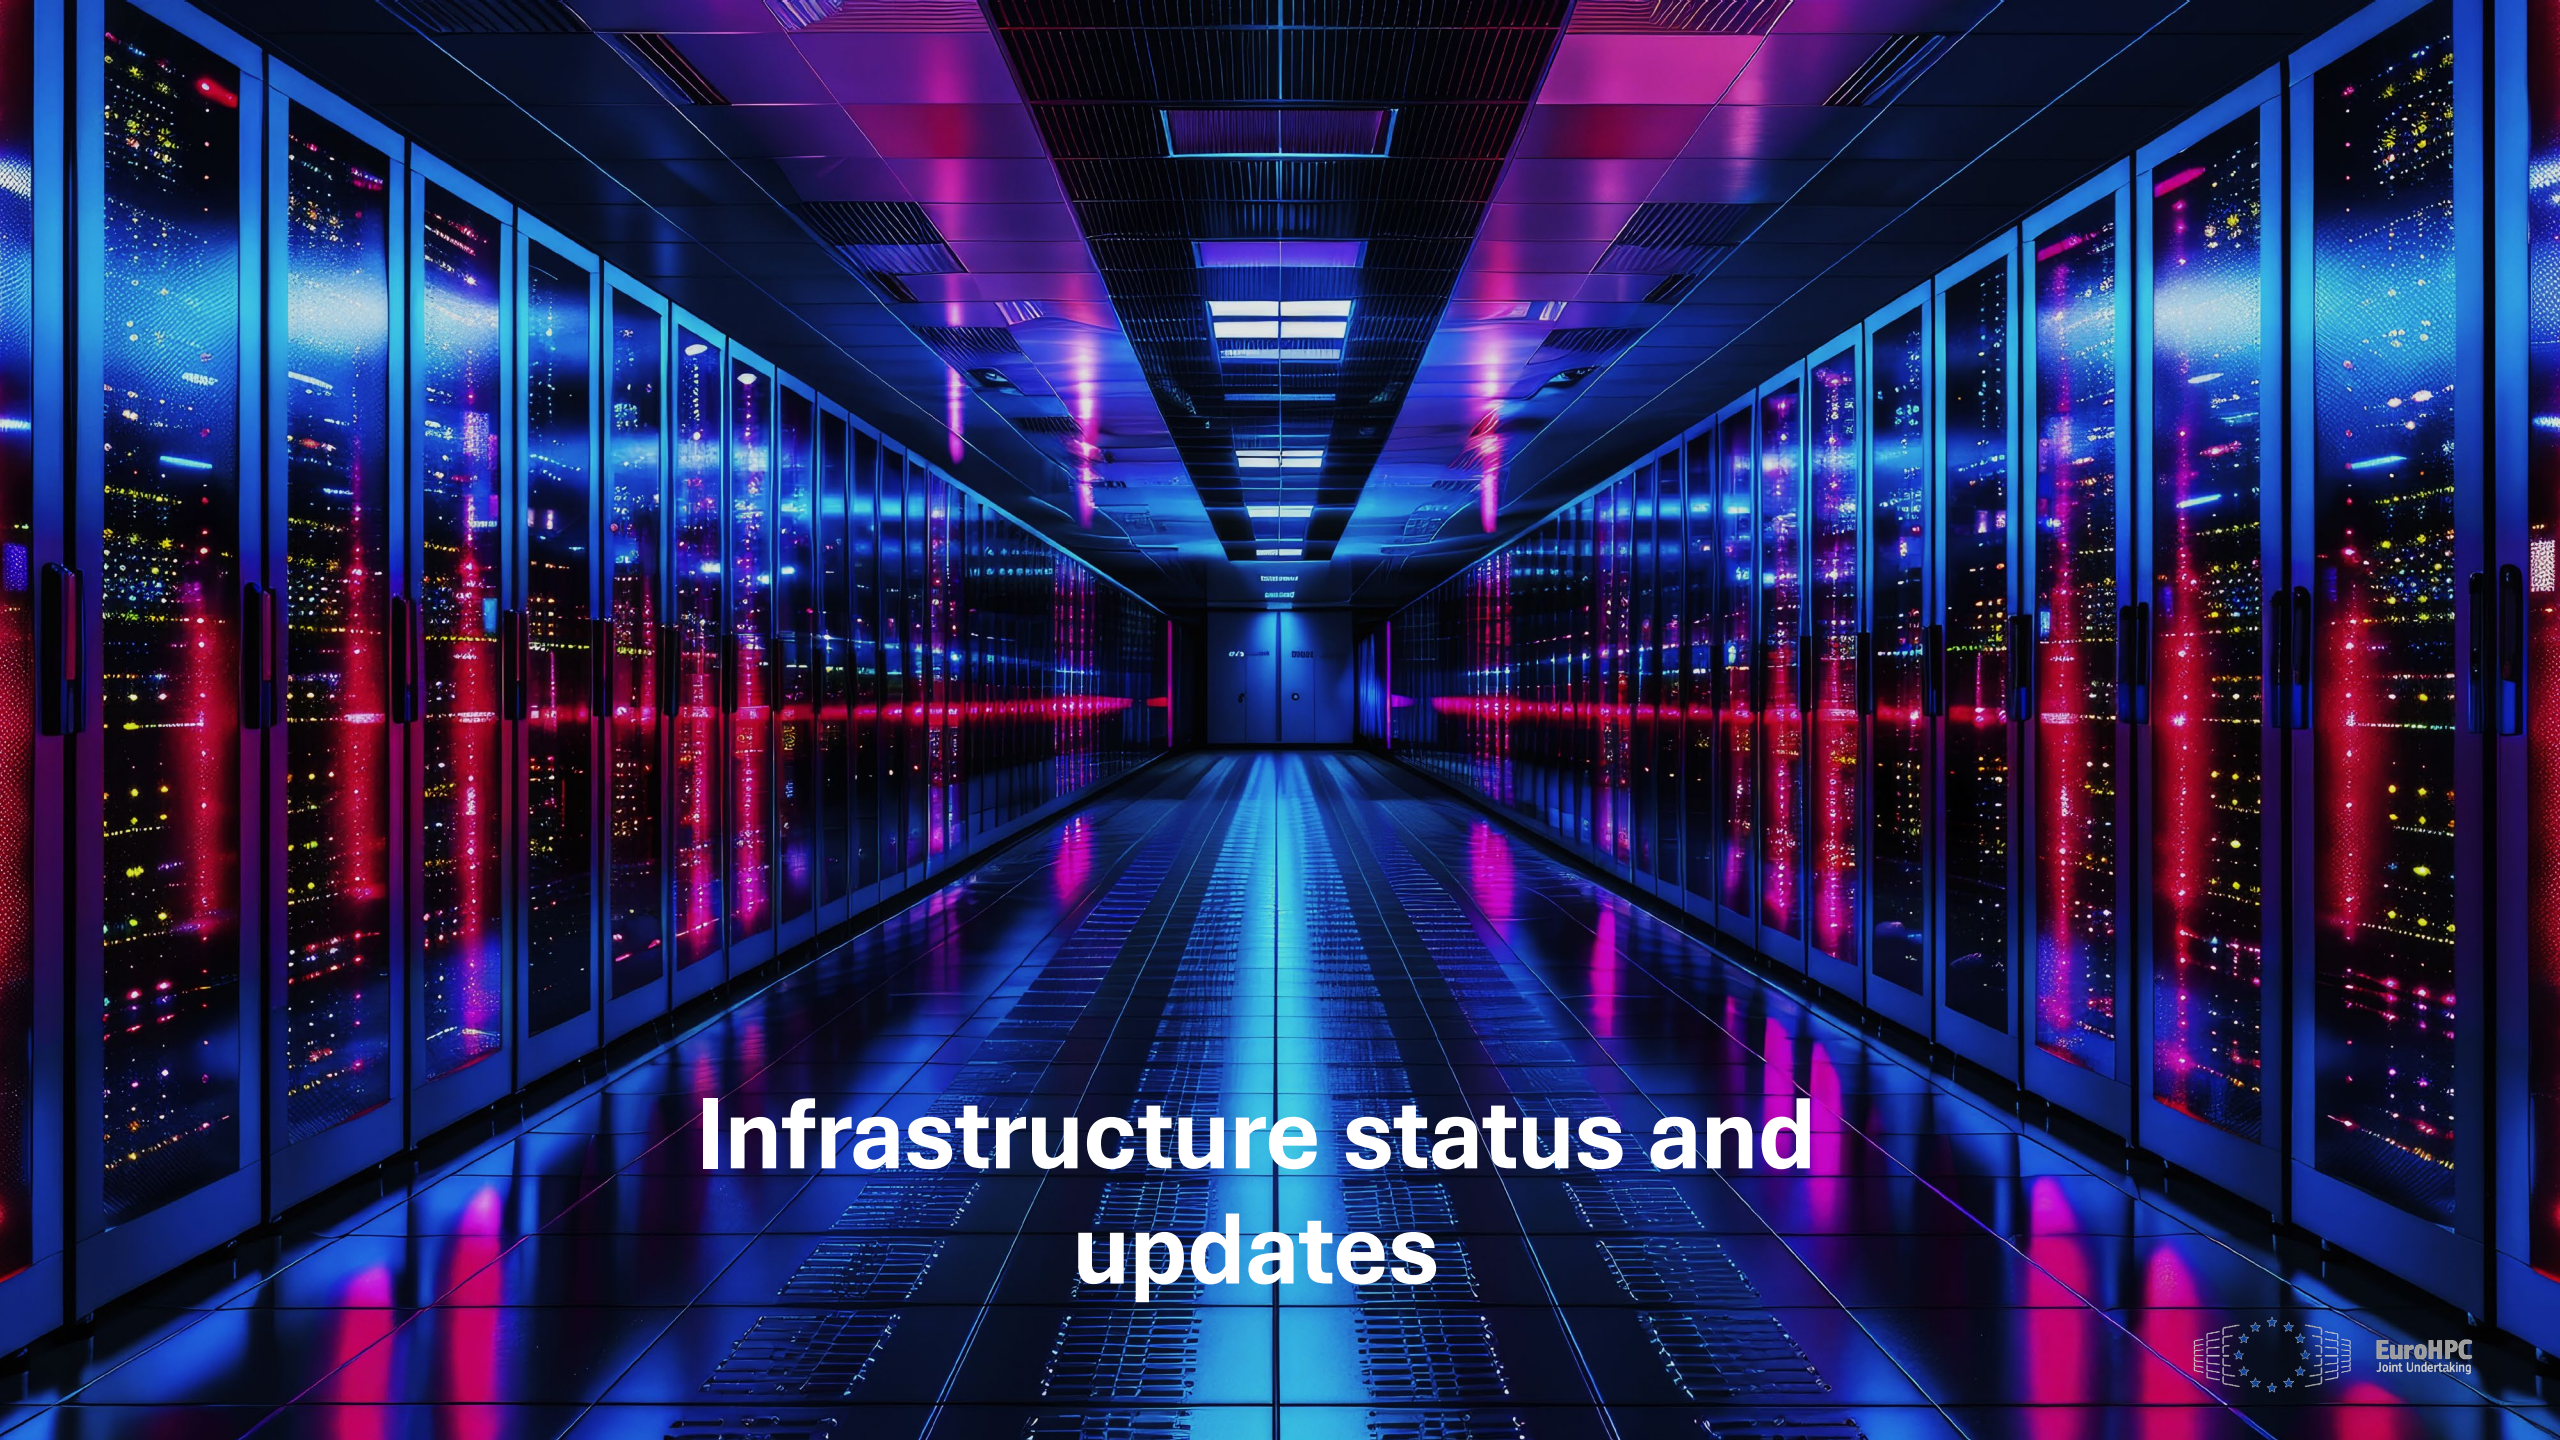

Infrastructure status and updates

EuroHPC
Joint Undertaking

The EuroHPC Supercomputing Ecosystem

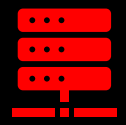

EXASCALE

PRE-EXASCALE

PETASCALE

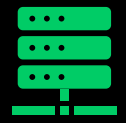

MID-RANGE

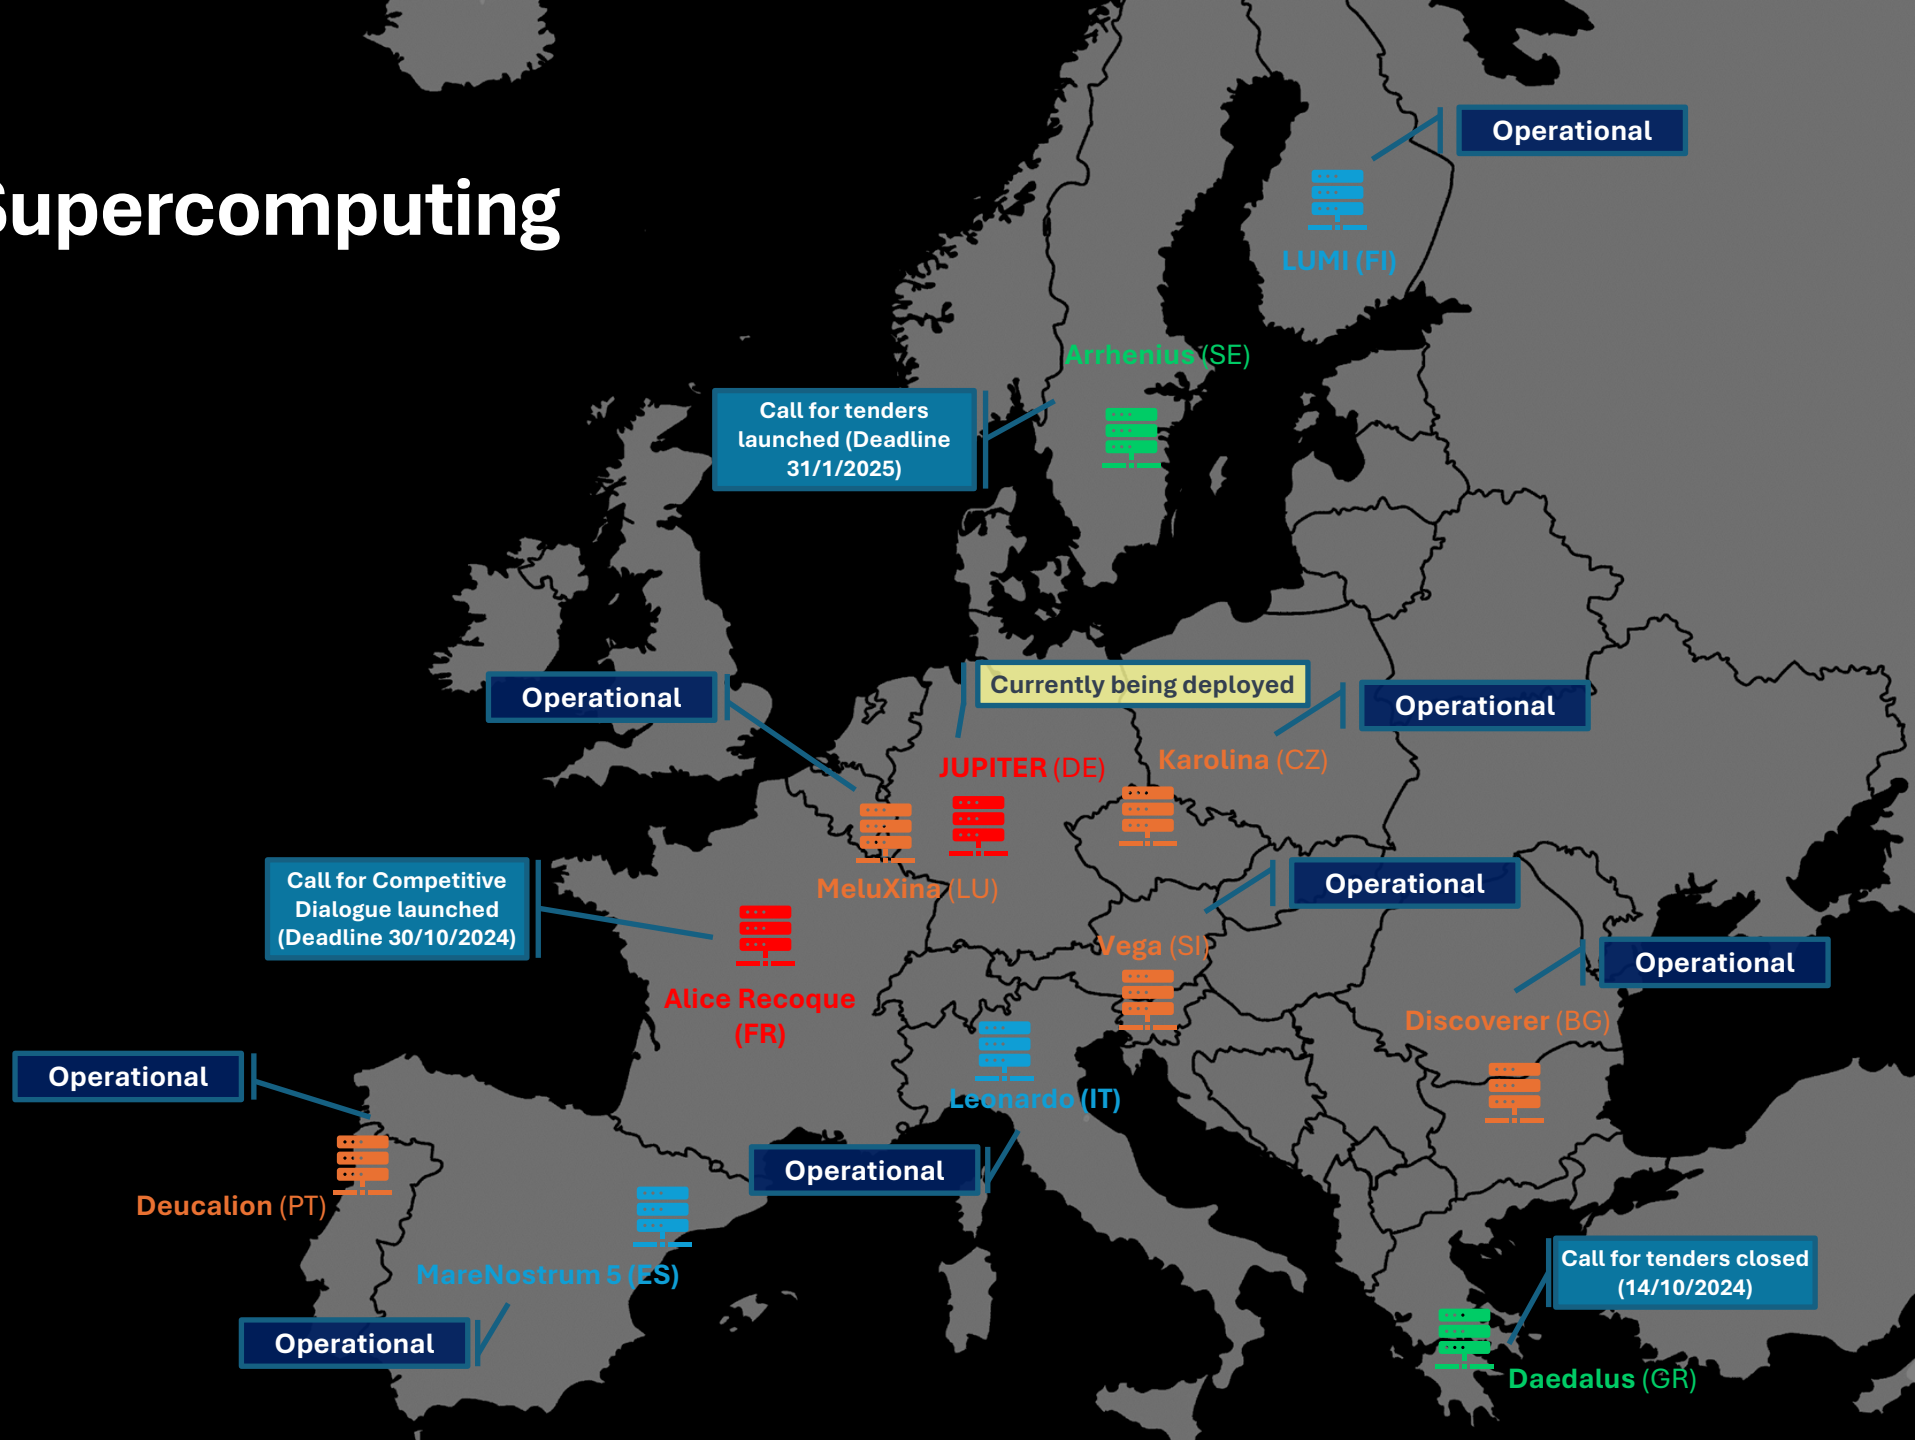

EuroHPC Joint Undertaking

Infrastructure updates

PRE-EXASCALE AND PETASCALE SYSTEMS

LISA – UPGRADE OF LEONARDO

Call launched on 18 September 2024

Closing date on 8 November 2024

Project start in March 2025

- The targeted system architecture is designed to address new evolving user needs involving AI workloads in the user workflows
- In conjunction with the HPC capacity of Leonardo, LISA will offer an AI-optimised partition, complementing the computing service portfolio of the whole infrastructure

DISCOVERER+

- Hardware deliveries expected mid-October 2024
- GPU installation planned in November 2024

Infrastructure updates

EXASCALE SYSTEMS

JUPITER

- First deliveries of the modular data centre (MDC) on site
- JUPITER Research and Early Access Program (JUREAP) ongoing
- JEDI (JUPITER Exascale Development Instrument) installed and benchmarked in May - #1 in Green500 – To be offered through the Early Access Program

ALICE RECOQUE

Call for Participation to the Competitive Dialogue published on 9 September 2024

Deadline for submissions: 30 October 2024

- Modular system targeting to support traditional HPC, AI training workloads, combination of HPC with AI inference workflows

Infrastructure updates

MID-RANGE SYSTEMS

DAEDALUS

Call for tenders closed 14 October 2024

- Modular system (accelerated + CPU partitions)
- Target installation date Q4 2025

ARRHENIUS

Call for tenders open until 31 January 2025

The system is divided into several specialized capabilities, each with dedicated storage to optimize performance for different types of workloads:

- **HPC CPU Module**
- **HPC GPU Module**
- **Sensitive Data Capability**
- **Persistent Compute and Data Services Capability (PCD)**

Access to the EuroHPC JU infrastructure

EuroHPC
Joint Undertaking

Access modes

OVERVIEW

- ✓ Allocations for 12 months
- ✓ Predefined minimum resource request and overall offer per cut-off
- ✓ 2 cut-offs per year

Extreme-Scale Access

For high-impact and high gain innovative research applications, with very large compute time, data storage and support needs.

Regular Access

For research and public sector applications requiring large-scale resources or frequent access to substantial computing and storage resources.

- ✓ Allocations for 12 months
- ✓ Predefined resources per partition
- ✓ Bi-monthly cut offs

AI and Data-Intensive Applications Access

For industry, SMEs, startups, and public sector entities requiring access to supercomputing resources to perform artificial intelligence and data-intensive activities.

- ✓ Allocations for up to 12 months
- ✓ Predefined resources per partition
- ✓ Monthly cut-offs

Development Access

For researchers and developers requiring a small number of node hours to develop, test and optimize their applications prior to applying for access.

Benchmark Access

Allows researchers and application developers to test or benchmark their applications.

WHO IS ELIGIBLE?

Principal Investigators and Team Members affiliated with organizations located in countries associated to Horizon 2020

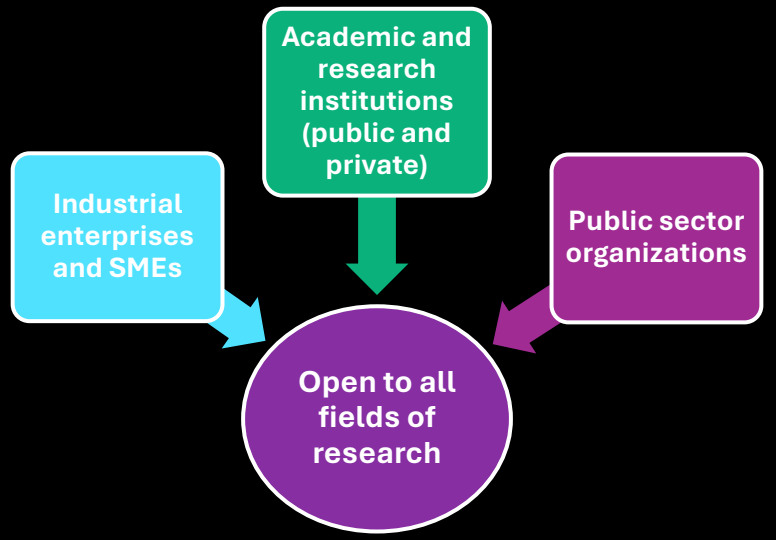

Access modes

CUT-OFFS TIMELINE

REGULAR ACCESS:

- December 2021
- March 2022
- July 2022
- November 2022
- March 2023
- July 2023
- November 2023
- March 2024
- September 2024 (under evaluation)

EXTREME SCALE ACCESS:

- December 2022
- May 2023
- October 2023
- April 2024
- October 2024 (under evaluation)

AI AND DATA INTENSIVE APPLICATIONS ACCESS :

- April 2024
- June 2024
- October 2024 (under evaluation)

UPCOMING CUT-OFFS IN 2024/2025:

EXTREME SCALE ACCESS – April 2025 (exact date TBD)

REGULAR ACCESS – March 2025 (exact date TBD)

AI AND DATA INTENSIVE APPLICATIONS ACCESS – 22 November 2024

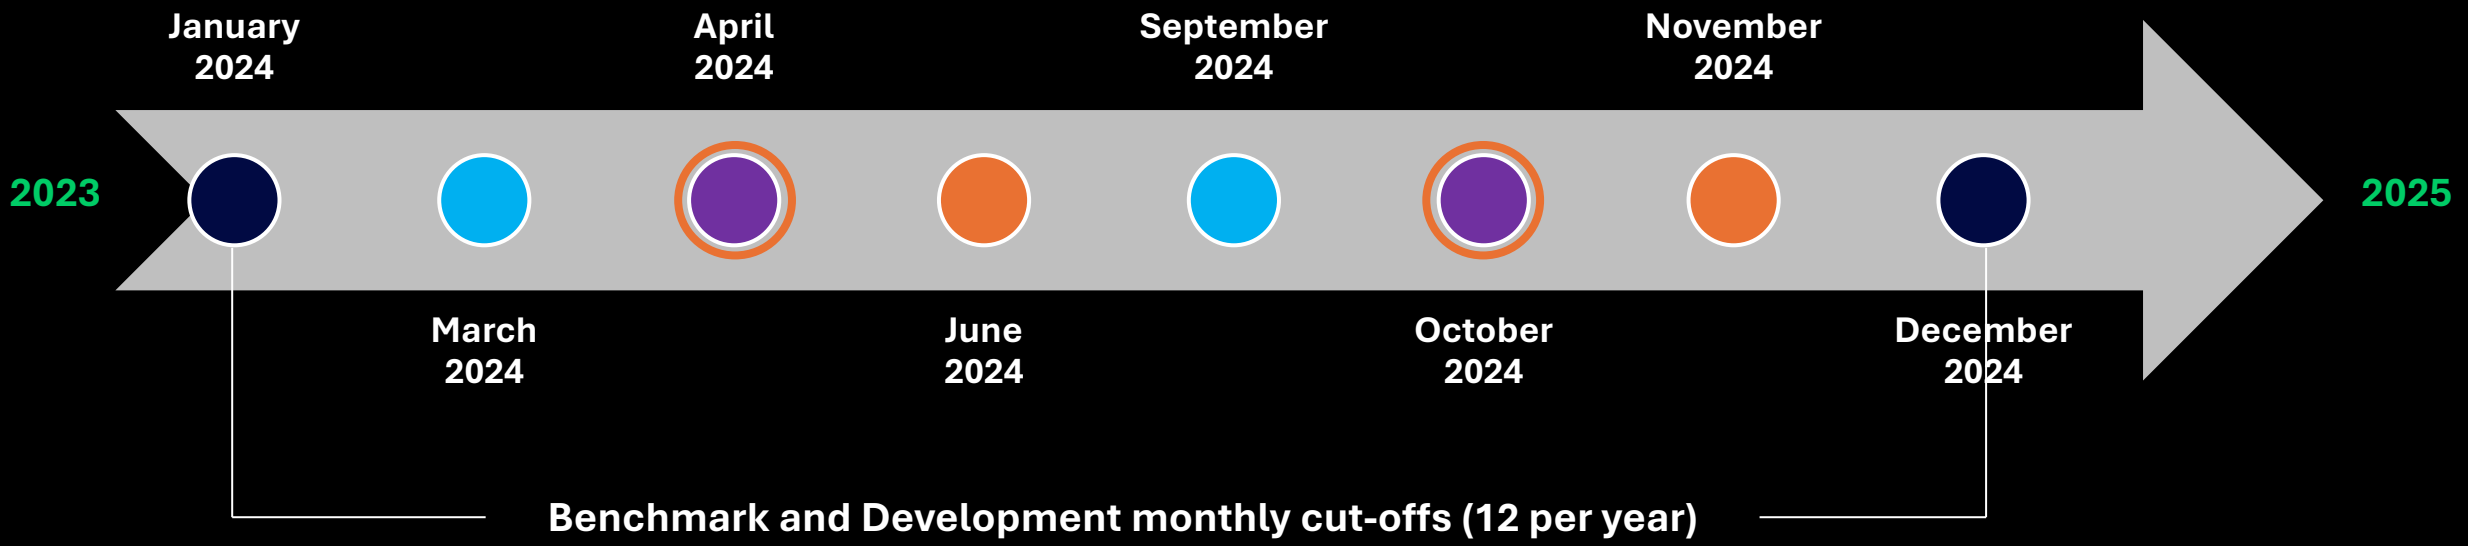

Peer-Review Process

EXTREME SCALE ACCESS

Access available on pre-exascale systems

Evaluation criteria:

Excellence
Innovation and Impact
Quality and Efficiency of the Implementation

Scoring system:

- Grade **0-5** per criterium
- Minimum grade per criterium – **3**
- Overall grade sum **0-15**
- Overall grade sum minimum - **10**

ACCESS TRACKS:

SCIENTIFIC

INDUSTRY

PUBLIC ADMINISTRATION

Peer-Review Process

REGULAR ACCESS

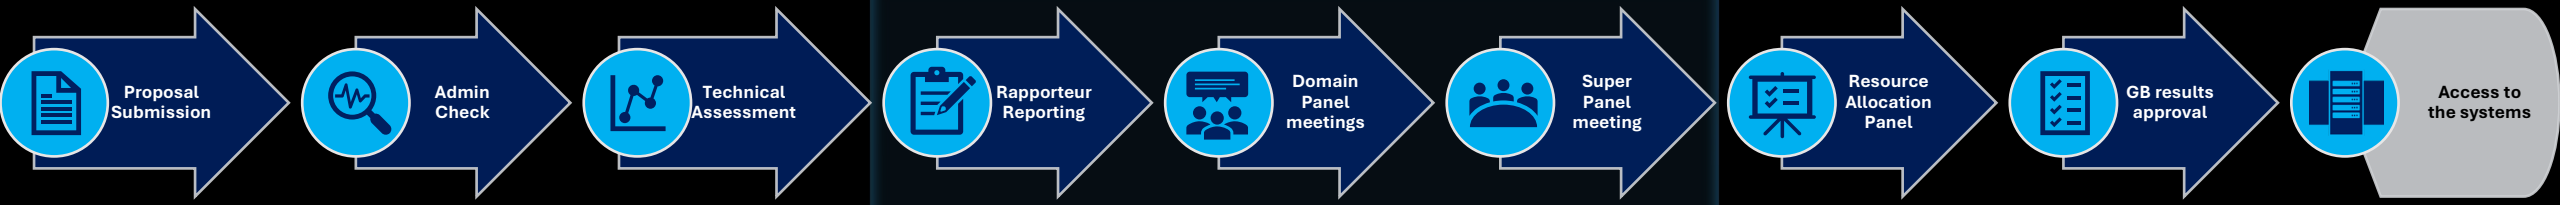

Access available on petascale and pre-exascale systems

Evaluation criteria:

Excellence
Innovation and Impact
Quality and Efficiency of the Implementation

Scoring system:

- Grade **0-5** per criterium
- Minimum grade per criterium – **3**
- Overall grade sum **0-15**
- Overall grade sum minimum - **10**

ACCESS TRACKS:

SCIENTIFIC

INDUSTRY

PUBLIC ADMINISTRATION

Peer-Review Process

AI AND DATA INTENSIVE APPLICATIONS ACCESS

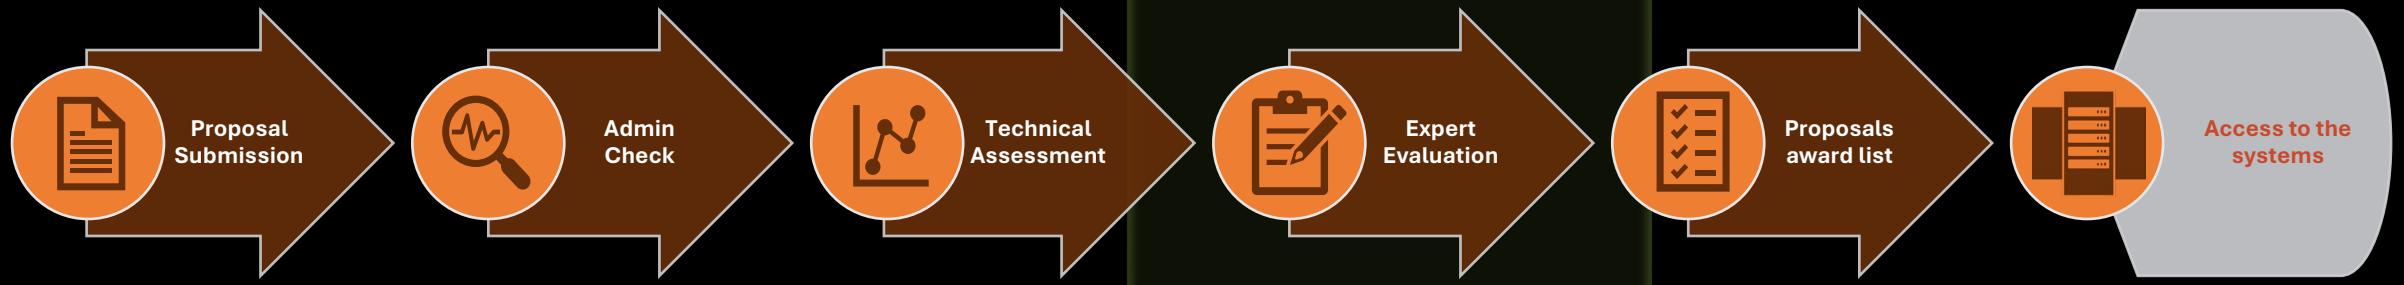

Evaluation criteria:

Excellence
Innovation and Impact
Quality and Efficiency of the Implementation

Scoring system:

- Grade **0-5** per criterium
- Minimum grade per criterium – **3**
- Overall grade sum **0-15**
- Overall grade sum minimum - **10**

Access available on GPU partitions of petascale and pre-exascale systems

Peer-Review Process

BENCHMARK AND DEVELEOPMENT ACCESS

Access available on petascale and pre-exascale systems

Evaluation process

INVOLVED ACTORS

THANK YOU!

ACCESS RESOURCE COMMITTEE ESTABLISHMENT IN 2024/2025

Evaluations of proposals' technical feasibility

Technical experts:

Computing centre representatives

Technical reviewers

Evaluations of proposals' scientific excellence, innovation and impact, quality and efficiency

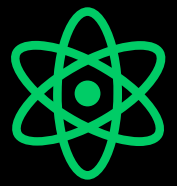

Scientific experts:

Committee Chairs

Domain Panel Chairs

Rapporteurs

External reviewers

Evaluation process

ADVICE FOR APPLICANTS

- Consult the EuroHPC JU website for updates
- Respect the cut-off dates and deadlines
- Use correct, up-to-date proposal templates
- Perform scalability tests on time on the preferred system
- For technical concerns contact the HPC centers
- Submit your Final Reports on time
- Take the Committee comments into consideration

Apply via:

https://access.eurohpc-ju.europa.eu/

Access calls statistics

EuroHPC
Joint Undertaking

Access calls statistics

OVERALL STATISTICS

AWARDED RESOURCES PER ACCESS MODE		
ACCESS CALL	**PROPOSALS AWARDED**	**NODE HOURS AWARDED**
EXTREME SCALE ACCESS (Dec 2022-Apr 2024)	75	63,113,698
REGULAR ACCESS (Dec 2021-Mar 2024)	189	25,698,394
AI AND DATA INTENSIVE APPLICATIONS ACCESS (Apr 2024-Jun 2024)	25	1,033,500
TOTAL	**289**	**89,845,592**

All calls for production activities - PI gender distribution - awarded projects

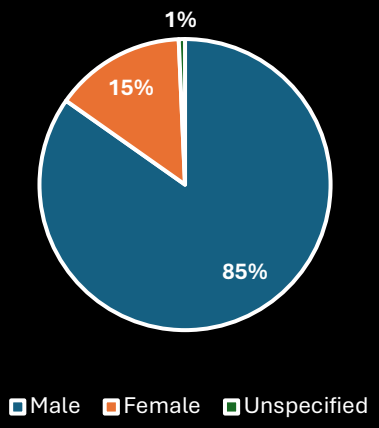

All calls for production activities - research domains distribution - awarded projects

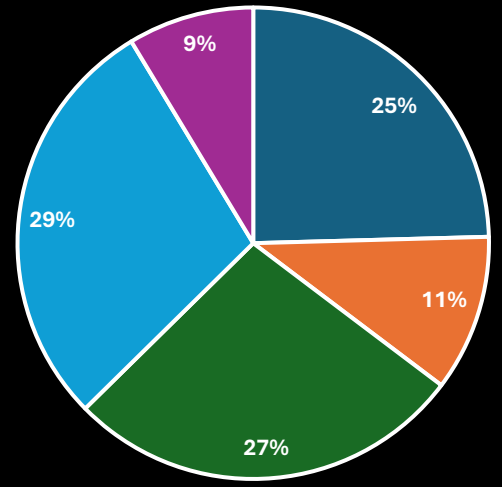

- Chemical Sciences and Materials, Solid State Physics
- Earth System Sciences & Environmental Studies
- Engineering, Mathematics and Computer Sciences
- Computational Physics: Universe Sciences, Fundamental Constituents of Matter
- Biochemistry, Bioinformatics, Life Sciences, Physiology and Medicine

Access calls statistics

OVERALL STATISTICS

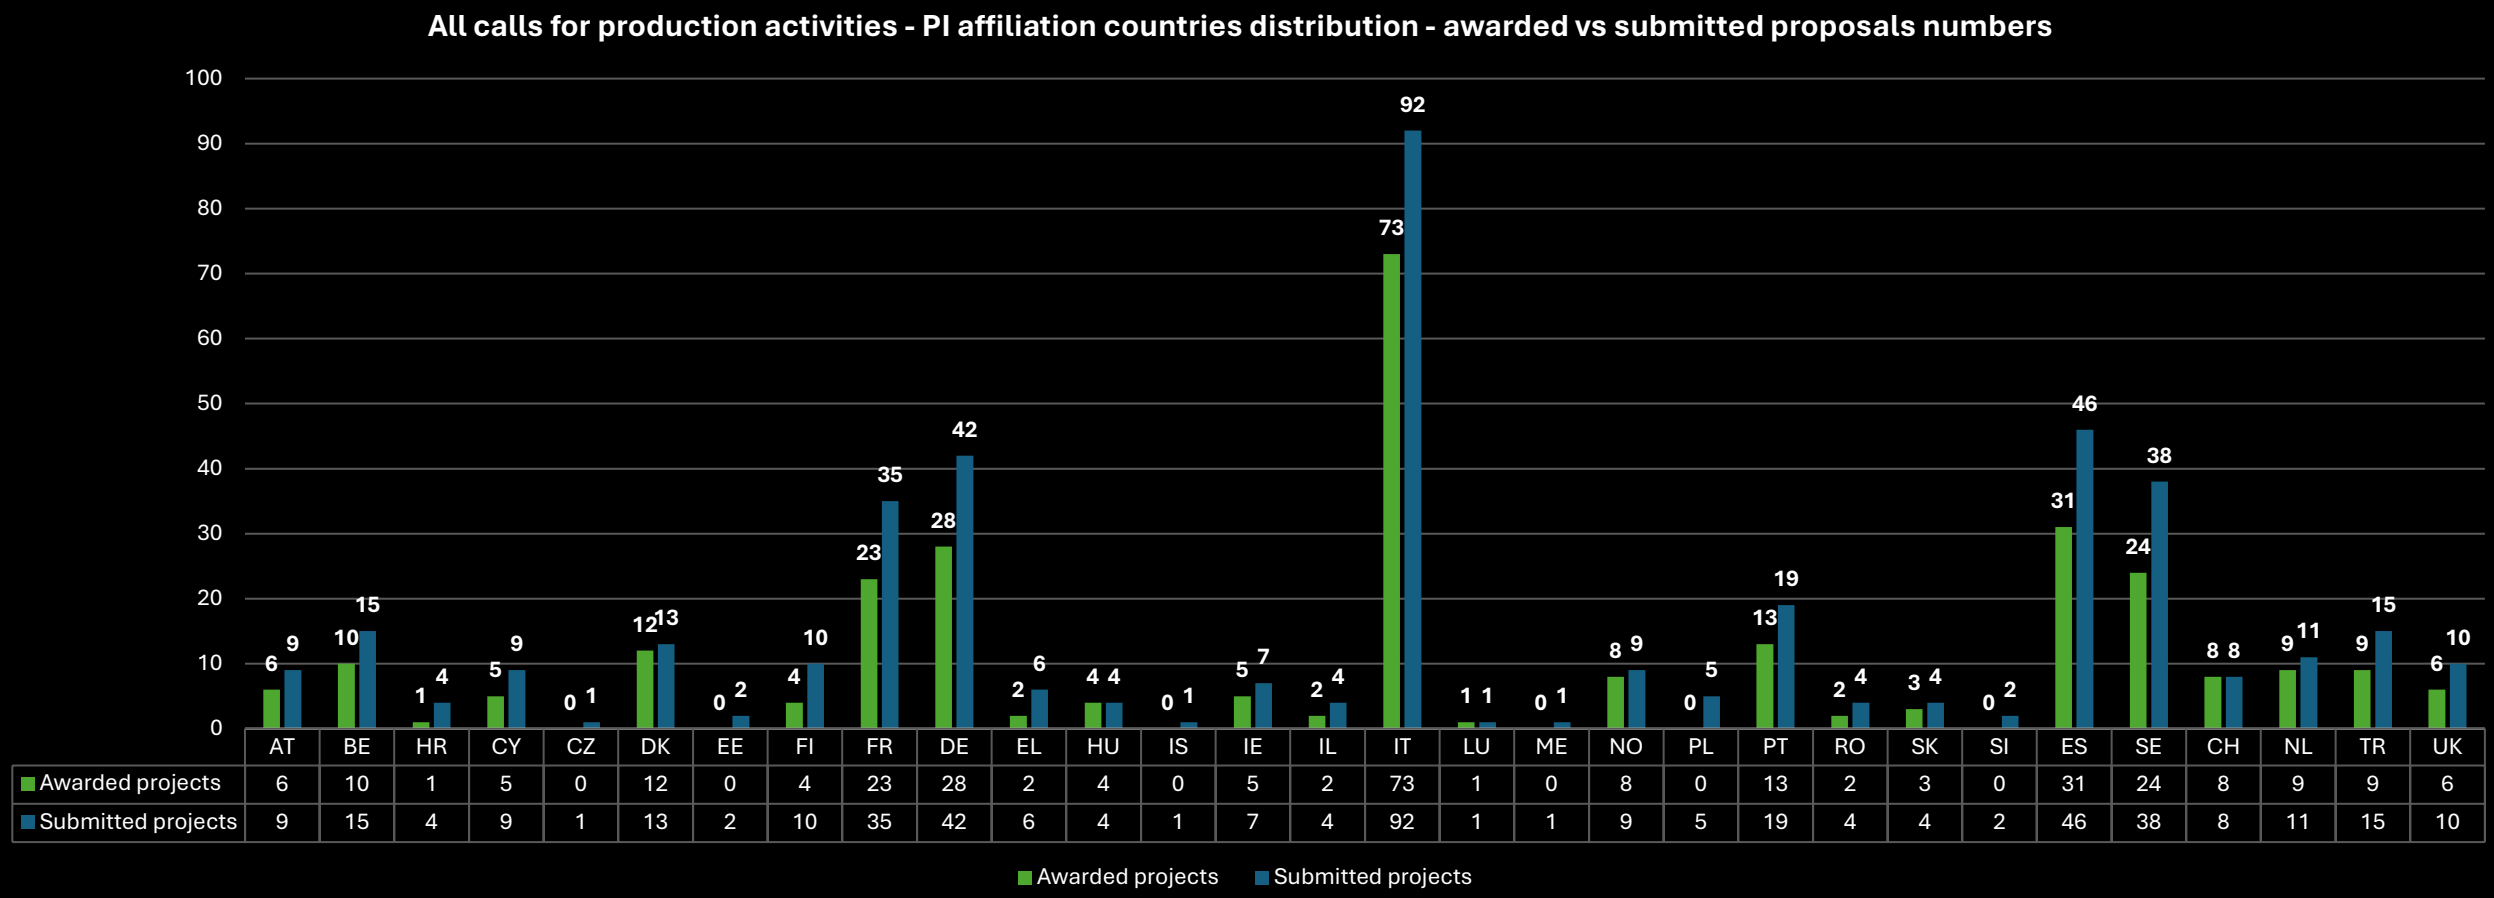

	AT	BE	HR	CY	CZ	DK	EE	FI	FR	DE	EL	HU	IS	IE	IL	IT	LU	ME	NO	PL	PT	RO	SK	SI	ES	SE	CH	NL	TR	UK
Awarded projects	6	10	1	5	0	12	0	4	23	28	2	4	0	5	2	73	1	0	8	0	13	2	3	0	31	24	8	9	9	6
Submitted projects	9	15	4	9	1	13	2	10	35	42	6	4	1	7	4	92	1	1	9	5	19	4	4	2	46	38	8	11	15	10

Access calls statistics

REGULAR ACCESS

Regular Access - Submitted vs administratively accepted vs awarded proposals (Dec 2021-Sep 2024)

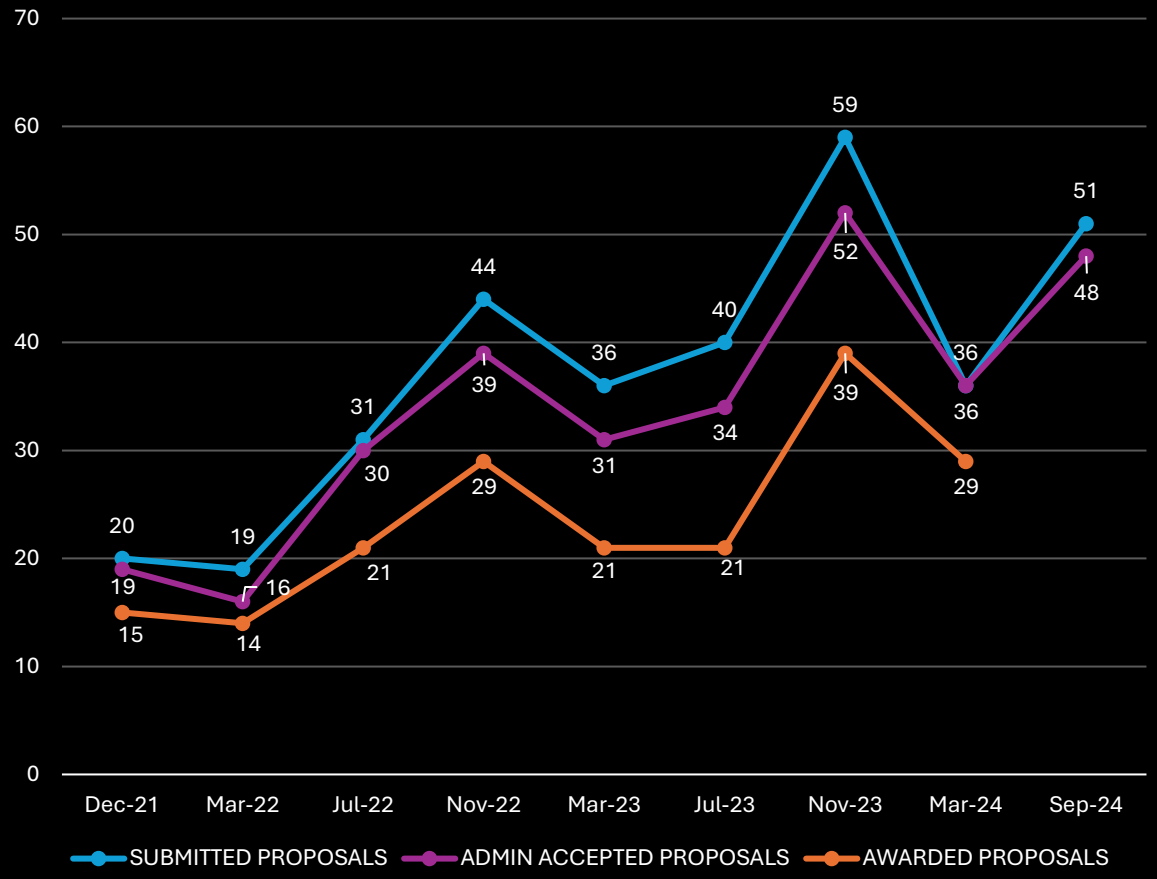

Regular Access - Research domains distribution of awarded proposals (Dec 2021-Mar 2024)

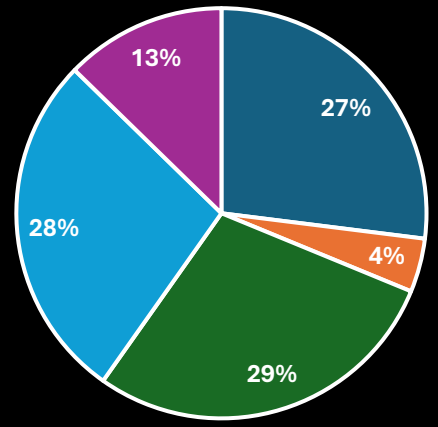

- Chemical Sciences and Materials, Solid State Physics
- Earth System Sciences & Environmental Studies
- Engineering, Mathematics and Computer Sciences
- Computational Physics: Universe Sciences, Fundamental Constituents of Matter
- Biochemistry, Bioinformatics, Life Sciences, Physiology and Medicine

Access calls statistics

REGULAR ACCESS

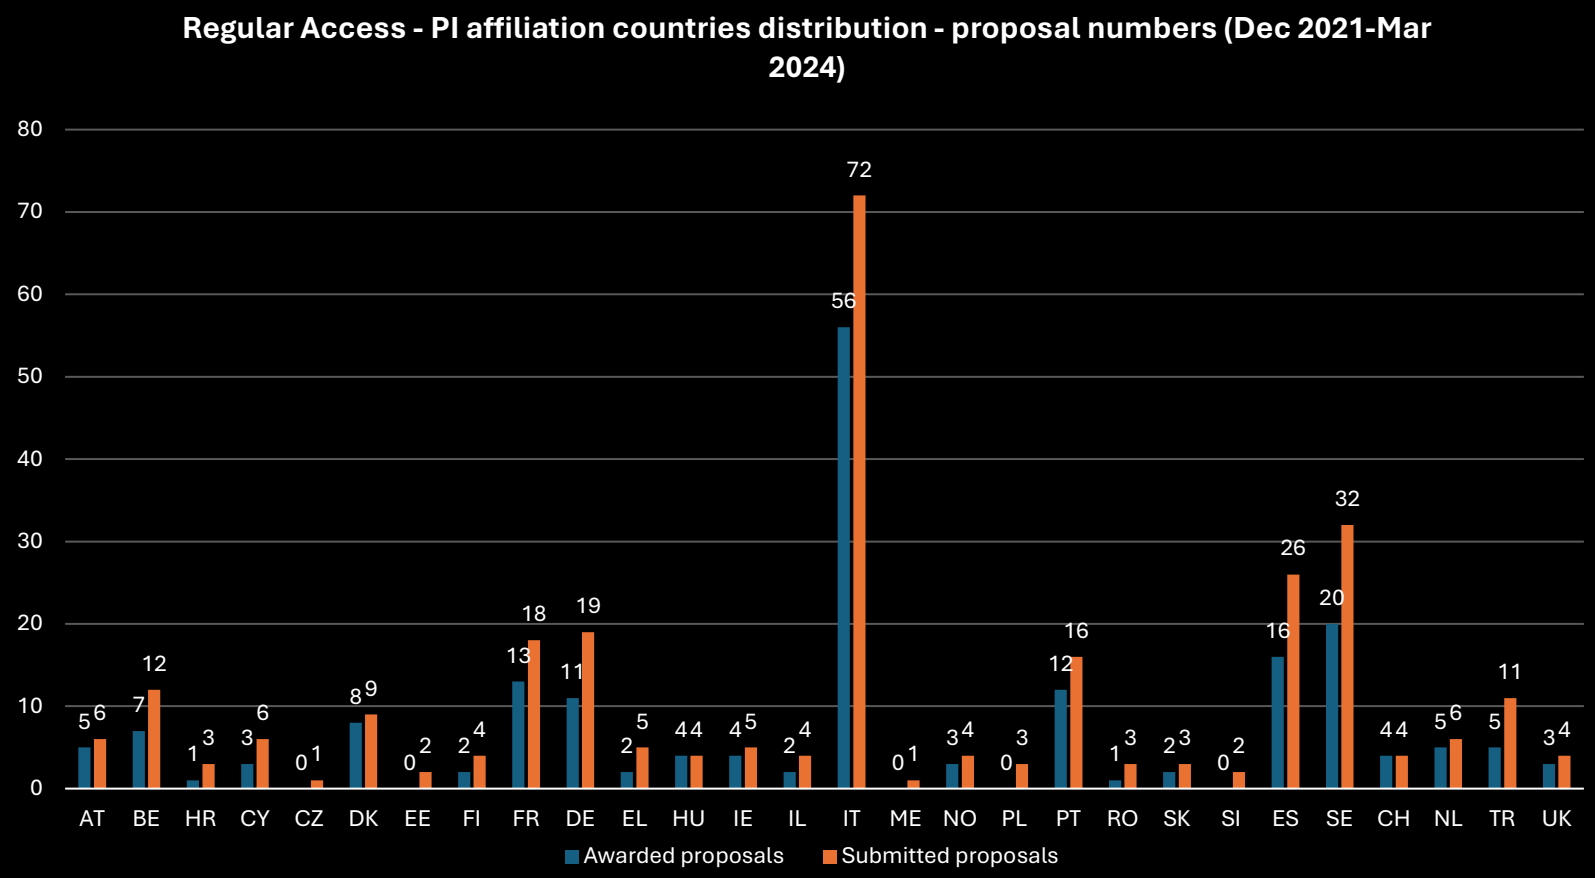

PROPOSAL NUMBERS PER COUNTRY - PI AFFILIATIONS			
COUNTRY	COUNTRY CODE	NO OF AWARDED PROJECTS	NO OF SUBMITTED PROJECTS
Austria	AT	5	6
Belgium	BE	7	12
Croatia	HR	1	3
Cyprus	CY	3	6
Czechia	CZ	0	1
Denmark	DK	8	9
Estonia	EE	0	2
Finland	FI	2	4
France	FR	13	18
Germany	DE	11	19
Greece	EL	2	5
Hungary	HU	4	4
Ireland	IE	4	5
Israel	IL	2	4
Italy	IT	56	72
Montenegro	ME	0	1
Norway	NO	3	4
Poland	PL	0	3
Portugal	PT	12	16
Romaina	RO	1	3
Slovakia	SK	2	3
Slovenia	SI	0	2
Spain	ES	16	26
Sweden	SE	20	32
Switzerland	CH	4	4
The Netherlands	NL	5	6
Türkiye	TR	5	11
United Kingdom	UK	3	4
TOTAL		**189**	**285**

Access calls statistics

REGULAR ACCESS

AWARDED RESOURCES (NODE HOURS) ACROSS ALL CUT-OFFS													
Cut-offs	Vega CPU	Vega GPU	MeluXina CPU	MeluXina GPU	Karolina CPU	Karolina GPU	Discoverer CPU	LUMI-C	LUMI-G	Leonardo Booster	Leonardo DCGP	MareNostrum5 GPP	MareNostrum5 ACC
Dec-21	328,125	31,923	0	173,325	177,344	0	156,250	1,865,234	0	0	0	0	0
Mar-22	882,160	7,813	0	147,344	151,563	0	135,240	281,142	0	0	0	0	0
Jul-22	1,102,710	104,688	122,200	227,600	273,438	56,250	0	1,343,281	0	0	0	0	0
Nov-22	508,049	29,688	468,750	163,705	392,896	49,297	890,625	2,488,506	0	0	0	0	0
Mar-23	763,573	0	493,753	125,781	468,750	46,875	990,005	0	0	0	0	0	0
Jul-23	451,650	15,866	110,000	174,650	86,240	21,828	325,000	177,000	288,807	296,810	53,177	36,938	0
Nov-23	604,400	50,000	511,719	190,320	493,900	48,000	1,096,000	386,500	547,520	300,563	140,000	38,724	286,404
Mar-2024	619,804	0	358,303	105,000	316,075	0	194,531	252,986	386,000	630,000	337,800	228,000	90,000
Total	**5,260,471**	**239,978**	**2,064,725**	**1,307,725**	**2,360,205**	**222,250**	**3,787,651**	**6,794,650**	**1,222,327**	**1,227,373**	**530,977**	**303,662**	**376,404**

HPC Centre	Core hours awarded	Node hours awarded
IZUM (SI)	704,057,247	5,500,447
IT4I (CZ)	330,554,171	2,582,454
SofiaTech (BG)	484,819,274	3,787,651
LuxProvide (LU)	385,800,563	3,372,450
CSC (FI)	921,970,288	8,016,977
CINECA (IT)	98,745,360	1,758,350
BSC (ES)	46,055,000	680,065
TOTAL	**2,972,001,903**	**25,698,394**

25 million node hours awarded via the Regular Access call

Access calls statistics

EXTREME SCALE ACCESS

Extreme Scale Access - Submitted vs administratively accepted vs awarded proposals (Dec 2022-Oct 2024)

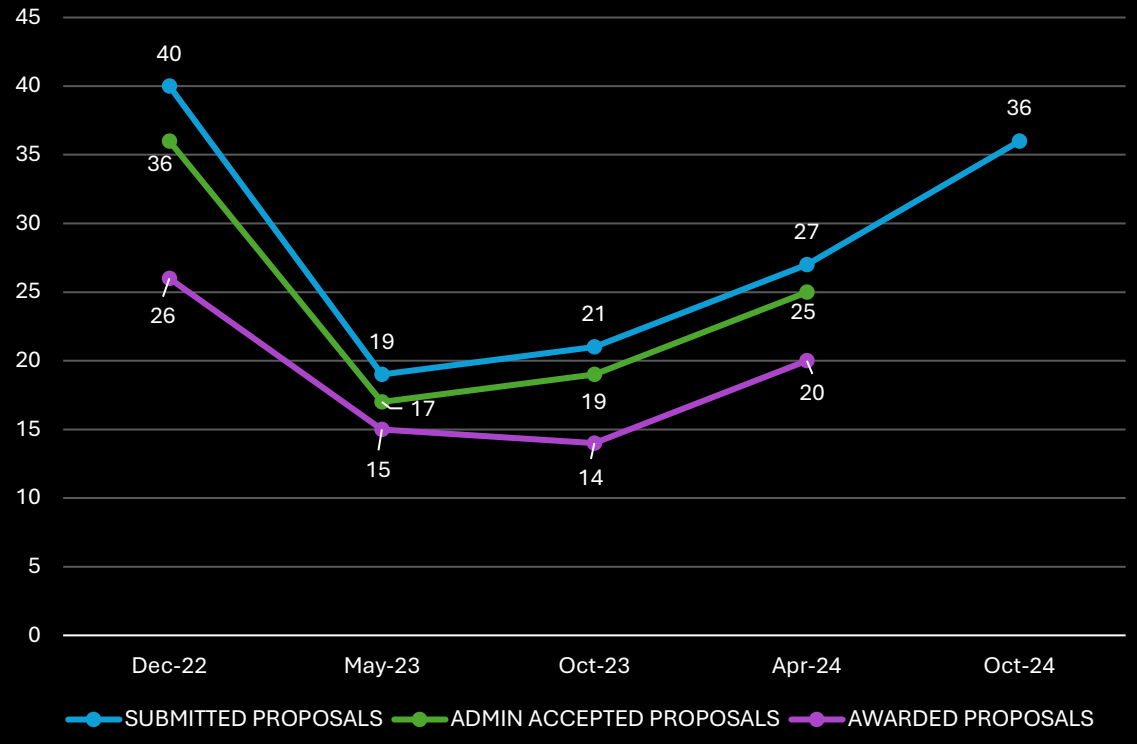

Extreme Scale Access - Research domains distribution of awarded proposals (Dec 2022-Apr 2024)

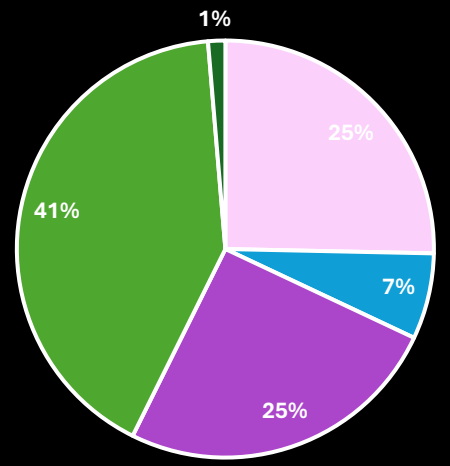

- Chemical Sciences and Materials, Solid State Physics
- Earth System Sciences & Environmental Studies
- Engineering, Mathematics and Computer Sciences
- Computational Physics: Universe Sciences, Fundamental Constituents of Matter
- Biochemistry, Bioinformatics, Life Sciences, Physiology and Medicine

Access calls statistics

EXTREME SCALE ACCESS

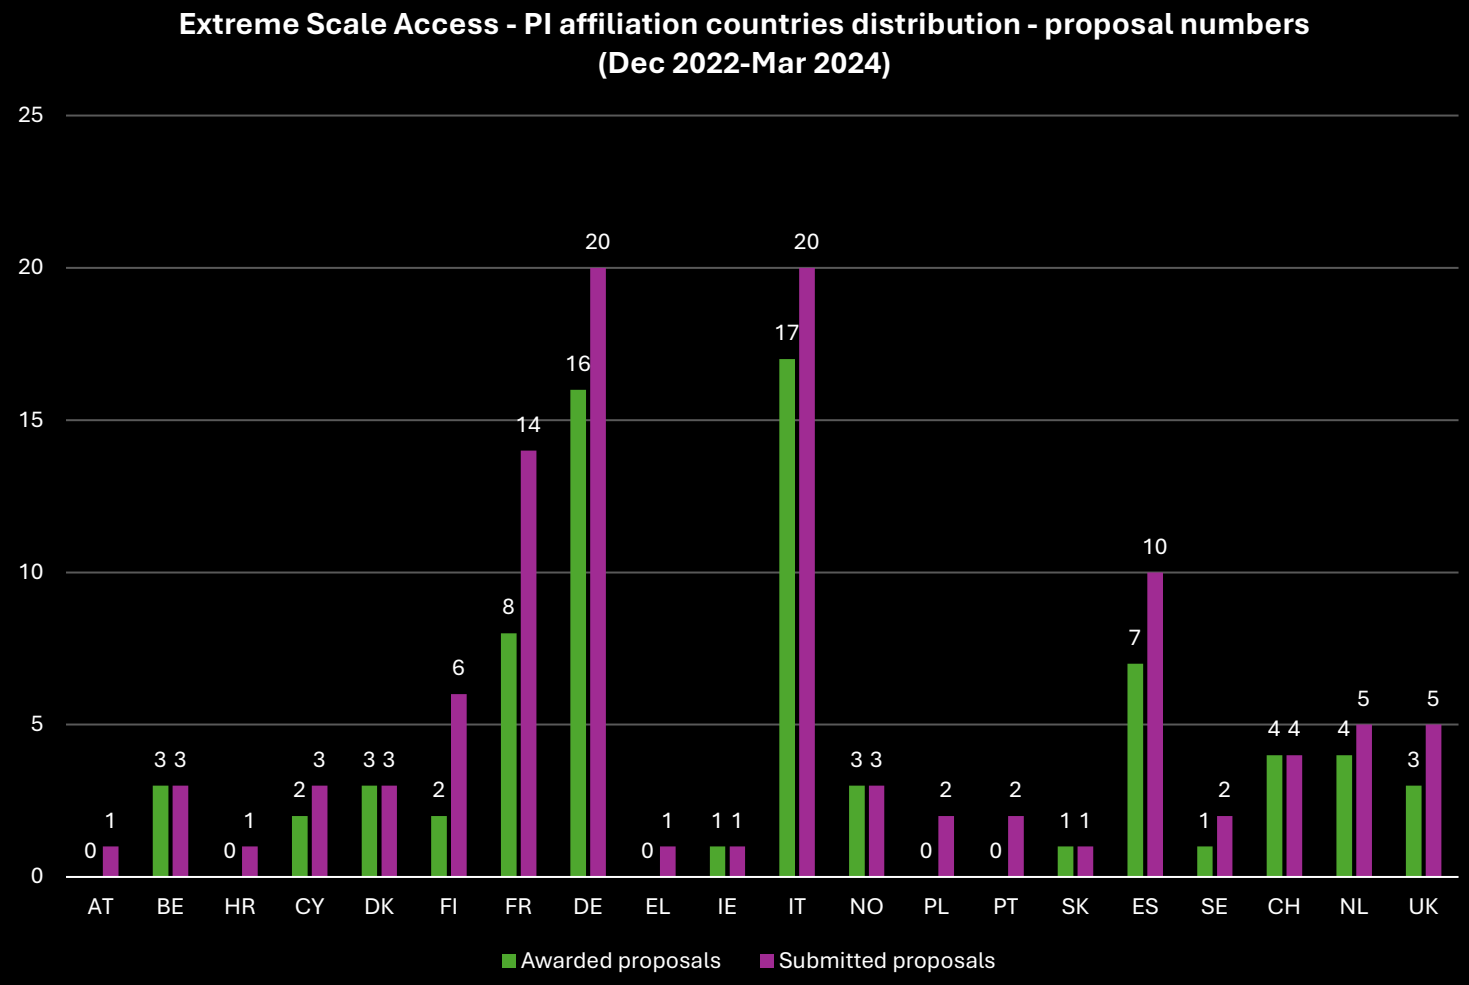

PROPOSAL NUMBERS PER COUNTRY – PI AFFILIATIONS			
COUNTRY	COUNTRY CODE	NO OF AWARDED PROJECTS	NO OF SUBMITTED PROJECTS
Austria	AT	0	1
Belgium	BE	3	3
Croatia	HR	0	1
Cyprus	CY	2	3
Denmark	DK	3	3
Finland	FI	2	6
France	FR	8	14
Germany	DE	16	20
Greece	EL	0	1
Ireland	IE	1	1
Italy	IT	17	20
Norway	NO	3	3
Poland	PL	0	2
Portugal	PT	0	2
Slovakia	SK	1	1
Spain	ES	7	10
Sweden	SE	1	2
Switzerland	CH	4	4
The Netherlands	NL	4	5
United Kingdom	UK	3	5
TOTAL		**75**	**107**

Access calls statistics

EXTREME SCALE ACCESS

AWARDED RESOURCES ACROSS ALL CUT-OFFS IN CORE AND NODE HOURS												
Cut-offs	CORE HOURS						NODE HOURS					
	Leonardo DCGP	Leonardo Booster	LUMI-C	LUMI-G	MareNostrum 5 GPP	MareNostrum 5 ACC	Leonardo DCGP	Leonardo Booster	LUMI-C	LUMI-G	MareNostrum 5 GPP	MareNostrum 5 ACC
Dec 2022	0	200,000,000	826,700,000	689,000,000	0	0	0	6,250,000	6,458,594	10,765,625	0	0
May 2023	0	144,659,008	436,667,904	505,024,000	0	38,063,648	0	4,520,594	3,411,468	7,891,000	0	1,189,489
Oct 2023	0	112,400,224	249,973,376	273,881,600	44,800,000	55,446,752		3,512,507	1,952,917	4,279,400	400,000	1,732,711
Apr 2024	35,840,000	98,255,712	89,895,424	140,582,016	432,320,000	19,200,000	320,000	3,070,491	702,308	2,196,594	3,860,000	600,000
Total	**35,840,000**	**555,314,944**	**1,603,236,704**	**1,608,487,616**	**477,120,000**	**112,710,400**	**320,000**	**17,353,592**	**12,525,287**	**25,132,619**	**4,260,000**	**3,522,200**

HPC Centre	Core hours awarded	Node hours awarded
CSC (FI)	3,211,724,320	37,657,906
CINECA (IT)	591,154,944	17,673,592
BSC (ES)	589,830,400	7,782,200
TOTAL	**4,392,709,664**	**63,113,698**

63 million node hours awarded via the Extreme Scale Access call

Access calls statistics

AI AND DATA INTENSIVE APPLICATIONS ACCESS

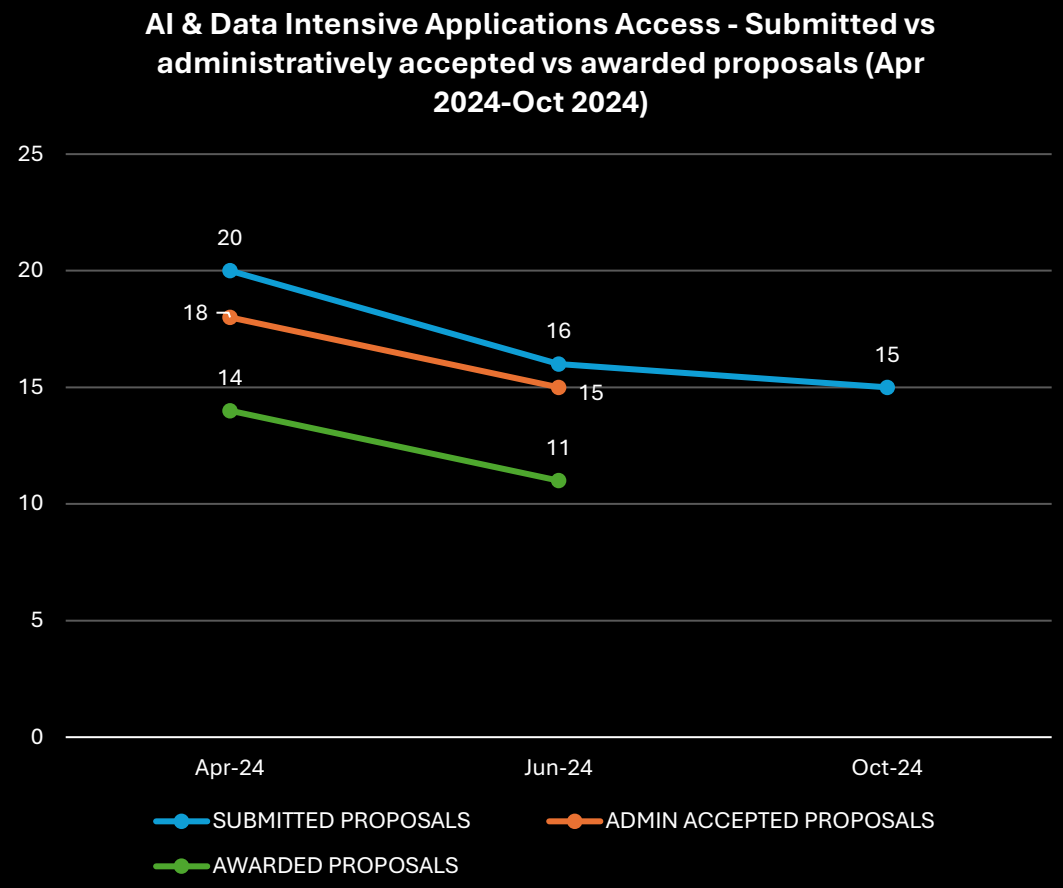

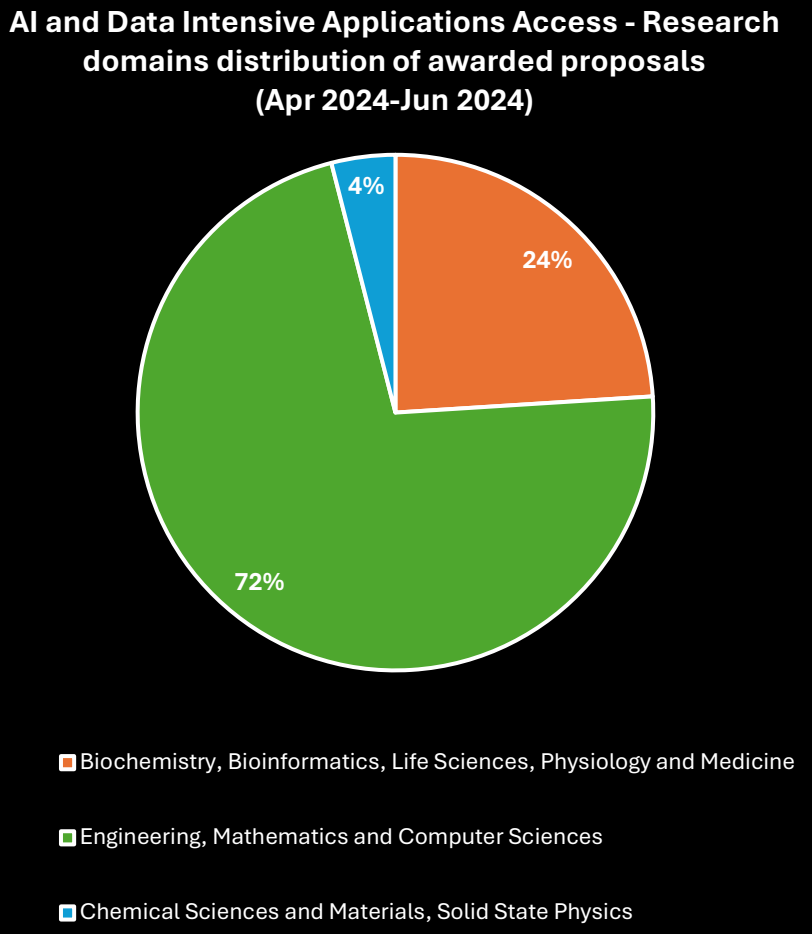

Access calls statistics

AI AND DATA INTENSIVE APPLICATIONS ACCESS

AI and Data Intensive Applications Access – Industry involvement – awarded proposal numbers (Apr 2024-Jun 2024)

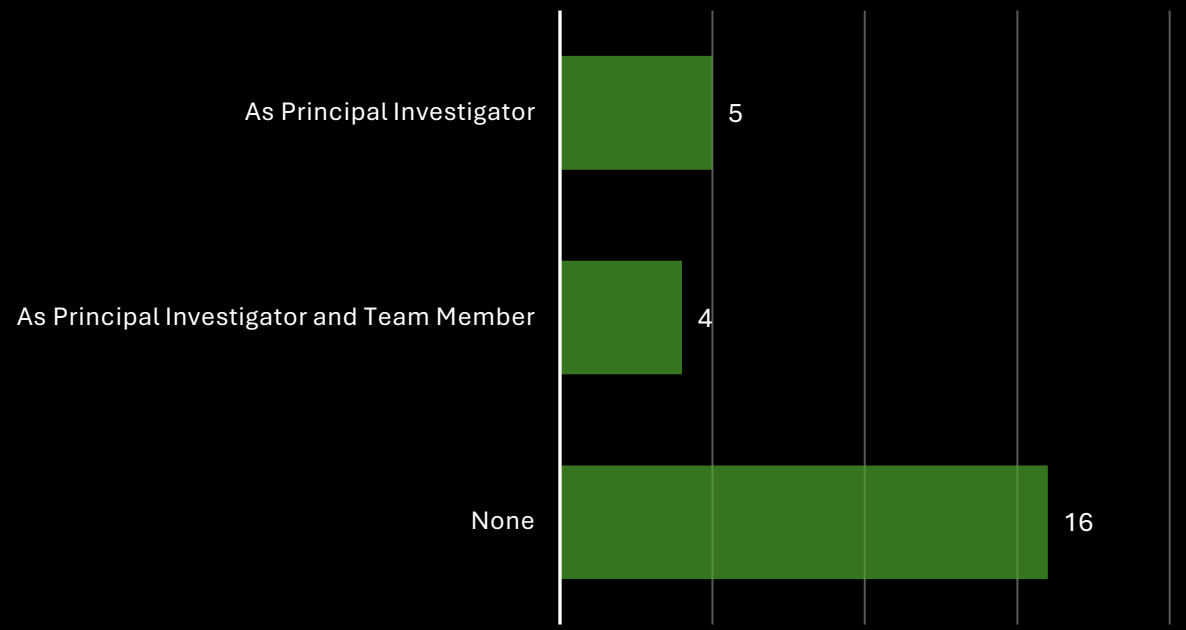

AI and Data Intensive Applications Access – AI technologies selected – submitted proposal numbers (Apr 2024-Oct 2024)

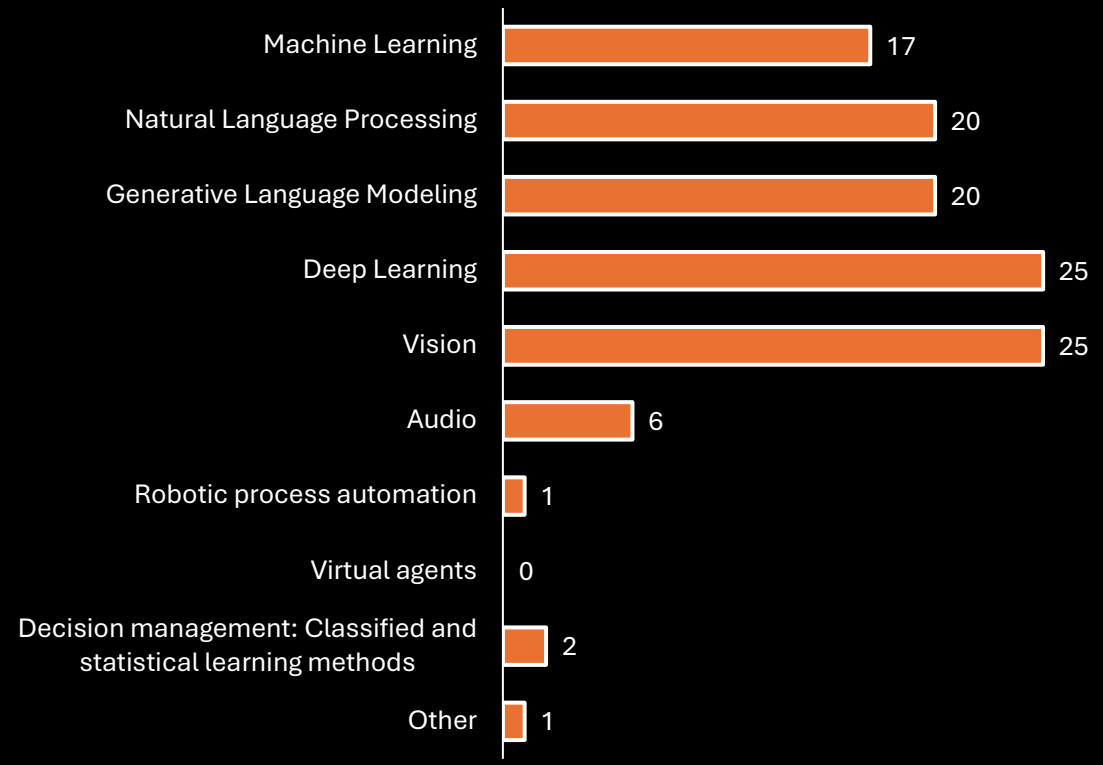

Access calls statistics

AI AND DATA INTENSIVE APPLICATIONS ACCESS

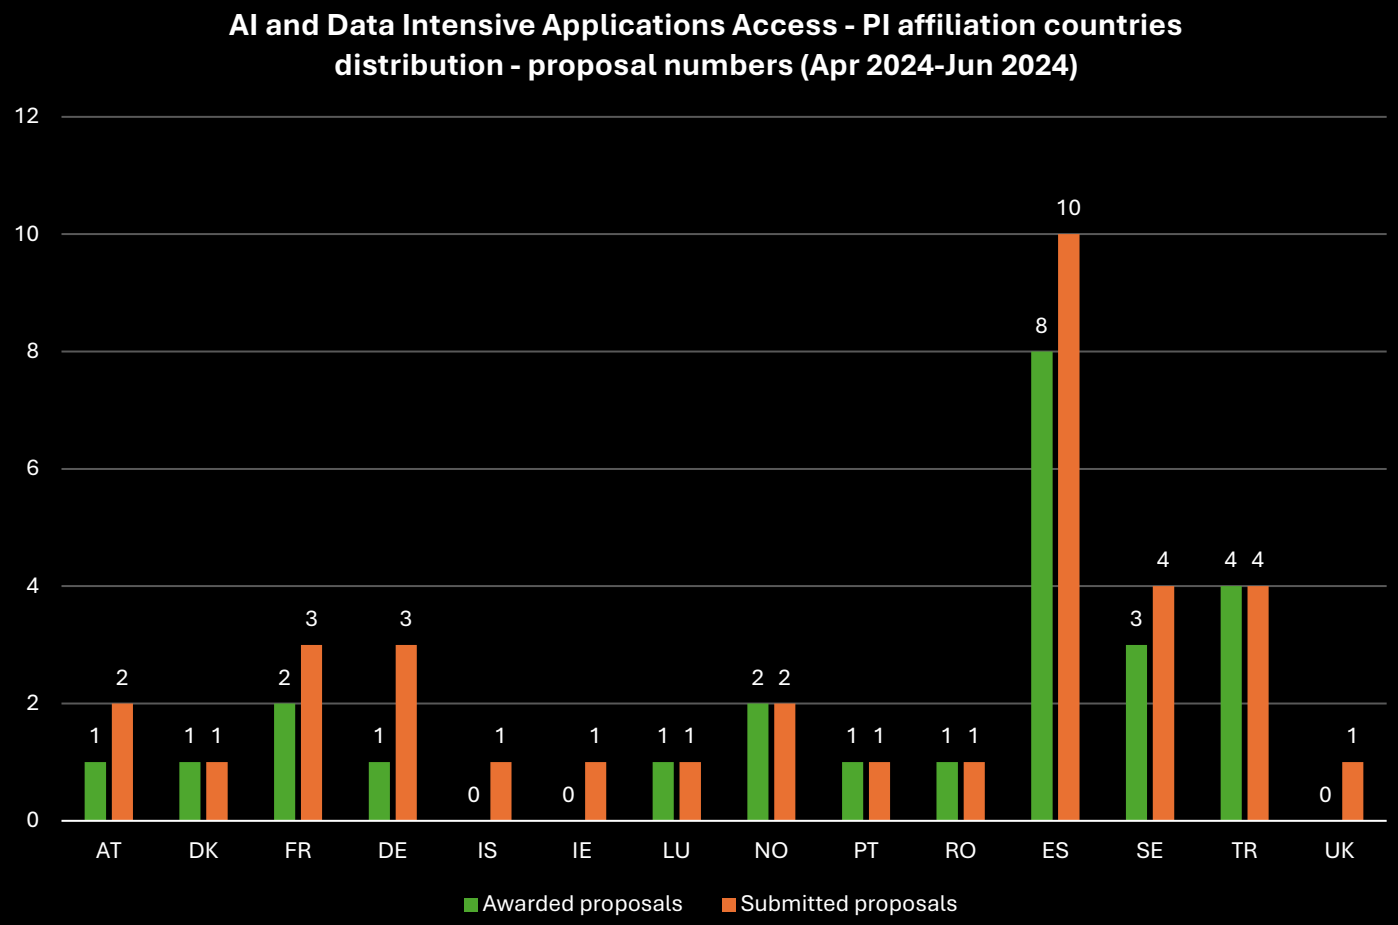

PROPOSAL NUMBERS PER COUNTRY - PI AFFILIATIONS			
COUNTRY	COUNTRY CODE	NO OF AWARDED PROJECTS	NO OF SUBMITTED PROJECTS
Austria	AT	1	2
Denmark	DK	1	1
France	FR	2	3
Germany	DE	1	3
Iceland	IS	0	1
Ireland	IE	0	1
Luxembourg	LU	1	1
Norway	NO	2	2
Portugal	PT	1	1
Romaina	RO	1	1
Spain	ES	8	10
Sweden	SE	3	4
Turkey	TR	4	4
United Kingdom	UK	0	1
TOTAL		**25**	**35**

Access calls statistics

AI AND DATA INTENSIVE APPLICATIONS ACCESS

AWARDED RESOURCES ACROSS ALL CUT-OFFS (NODE HOURS)							
Cut-off	**Vega GPU**	**MeluXina GPU**	**Karolina GPU**	**LUMI-G**	**Leonardo Booster**	**MareNostrum5 ACC**	**TOTAL**
April 2024	0	0	7,500	35,000	400,000	128,000	570,500
June 2024	0	0	0	35,000	300,000	128,000	463,000
TOTAL	**0**	**0**	**7,500**	**70,000**	**700,000**	**256,000**	**1,033,500**

1 million node hours awarded via the AI and Data Intensive Applications Access call

SYSTEM OFFERS (NODE HOURS)	
Vega GPU	7,100
MeluXina GPU	25,000
Karolina GPU	7,500
LUMI-G	351,455
Leonardo Booster	545,865
MareNostrum5 ACC	129,377
TOTAL	**1,065,918**

Access calls statistics

BENCHMARK AND DEVELOPMENT ACCESS

Benchmark and Development Call Statistics for the period of Jan - Sept 2024

ACCESS MODE	NUMBER OF SUBMITTED PROJECTS	REQUESTED NODE HOURS	NUMBER OF AWARDED PROJECTS	AWARDED IN % OF REQUESTED PROJECTS	AWARDED NODE HOURS	AWARDED NODE HOURS IN % OF REQUESTED NODE HOURS
BENCMARK ACCESS	188	720.000	158	84%	629.800	87%
DEVELOPMENT ACCESS	300	2.450.900	246	82%	2.195.500	90%
Total	488	3.170.900	404	83%	2.825.300	89%

Access calls statistics

BENCHMARK AND DEVELOPMENT ACCESS

Benchmark and Development Call Statistics for the period of Jan - Sept 2024

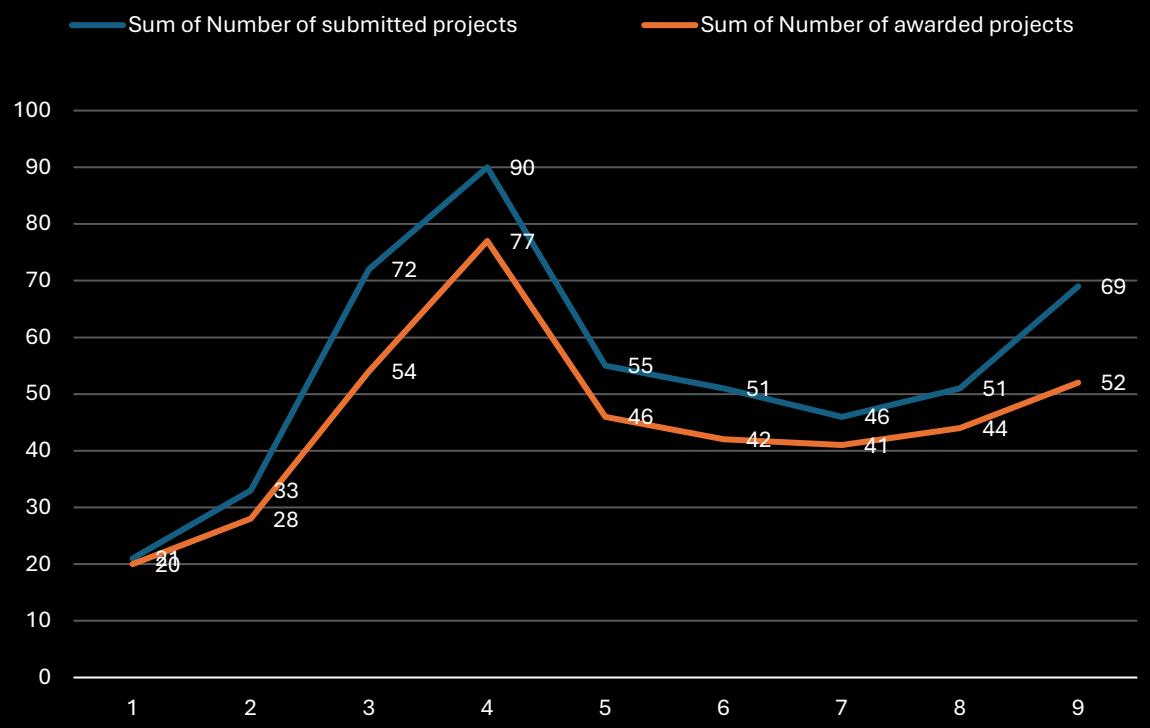

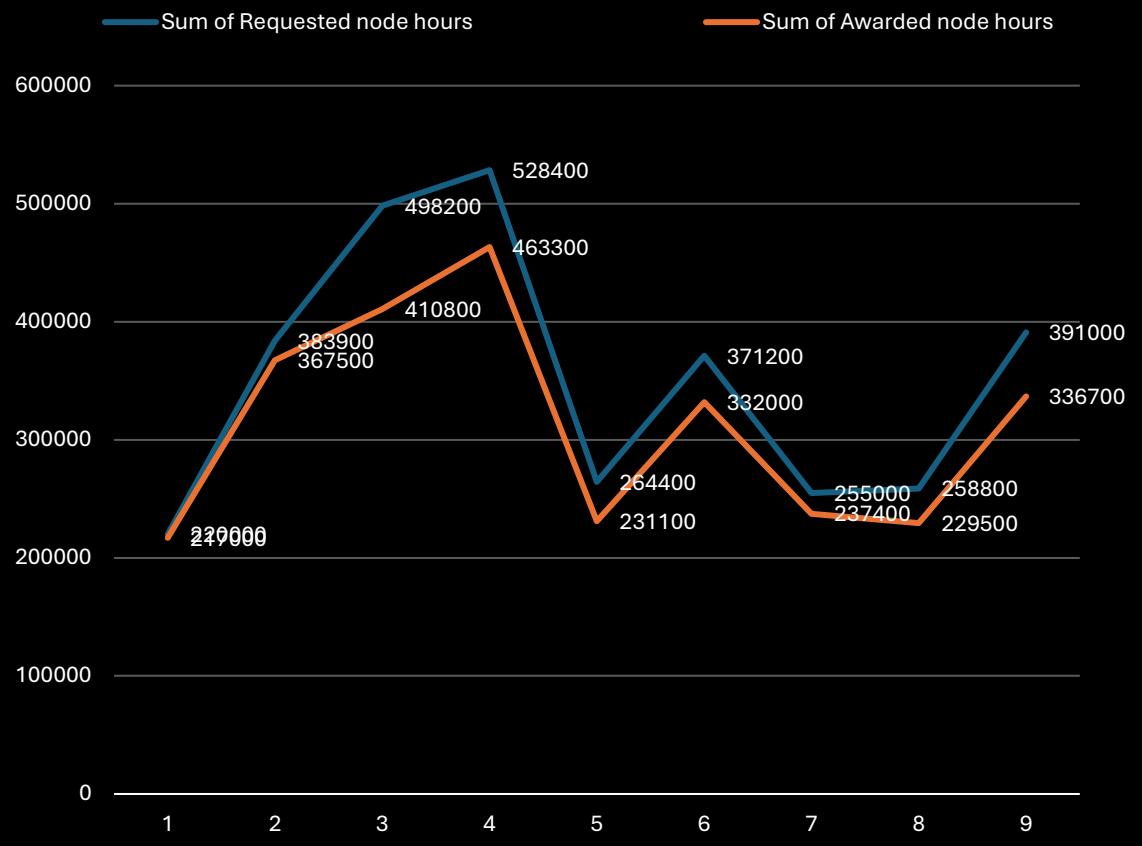

The Peer-Review Team

Klara Meštrović

Klara.Mestrovic@eurohpc-ju.europa.eu

Krishnakshi Bhuyan

Krishnakshi.Bhuyan@eurohpc-ju.europa.eu

Dora Marton

Dora.Marton@eurohpc-ju.europa.eu

Catarina Guerreiro

Catarina.Guerreiro@eurohpc-ju.europa.eu

access@eurohpc-ju.europa.eu

Thank you!

For more information, feel free to visit our website and social media:

/eurohpc-ju.europa.eu

/EuroHPC_JU

/eurohpc-ju

/eurohpc-ju

THE EUROPEAN HIGH PERFORMANCE COMPUTING JOINT UNDERTAKING

EuroHPC
Joint Undertaking

The Access Resource Committee

some comments from

Maria Paola Lombardo and Tobias Weinzierl

I- Access Resource Committee (ARC):

Generalities, set-up, operations

Involvement in different calls and domains

Different calls at a glance:

Regular Access

Domain-specific panels
10 pages
See max node hrs
4 months turnaround
Workhorse

Extreme-scale Access

One big panel
20 pages
See min node hrs
6 months turnaround
Champion runs

ARC members take part in Regular and Extreme Access Calls

Regular Access Calls :
Panel chair appoints panel (per domain)
Reviews (2/proposal) by panel members → panel discussion → domain ranking
Superpanel (i.e. all chairs) → panel discussion → eligible proposals, global ranking

Extreme Access Calls:
ARC chair and deputy appoint panel of rapporteurs (only one panel)
Rapporteurs appoint referees (3/proposal)
Reviews plus technical assessment → reply by the PI
Rapporteurs summarise and moderate reviews plus rebuttals internally
ARC discussion → grading of the proposals → eligible proposals, global ranking

Currently about 70 ARC members – our own involvement & expertise:

MpL : EuroHPC Expert – Computational Physics
TW : EuroHPC Expert – Computer Science

II- Best practices and suggestions to users

- As the majority of proposals are above threshold, our main challenge is ranking across different fields. Hence, a clear explanation of the expected impact on a given field is a strong asset of the proposal.
- Take the technical requirements seriously: if needed, use help from the Centers to produce scaling plots on the machines your are applying to, and optimize.
- Keep in mind we have to select projects that can exploit the resources in the best possible way: motivate your requests and explain why certain resources requested.
- Statements on possible “minimum acceptable allocation” (although not required) are helpful.
- We are evaluating proposals: Even if you are a leading expert in the field and/or if you are working on a well-known important problem, motivations and details are needed.

III- Frequently seen weaknesses

Technical assessment/feedback not worked in

I/O and data transfer crucial

HPC terminology (e.g. speedup) not properly used

Evolutionary character not clearly justified

Lack of methodological innovation

Continuation/evolutionary proposals without clear outcomes/dissemination of previous submission (plus self-plagiarism, added motivation … yet, they are often necessary)

Clear capability character of proposals

IV. For our referees

EuroHPC Joint Undertaking

External Referees play a crucial role in the Extreme Access Calls - We thank them!

For their work to be useful, please note: The panel does not average the grades from different referees. Rather, the panel seeks coherence, tries to understand differences, and to reach a consensus.

Hence:

– Try to be as objective as possible in your evaluation, highlighting strength and weaknesses of the proposal - unduly generous statements and/or generic negative statements are not useful.

–Strive for coherence between your comments and the marks. Not only incoherent marks may receive little consideration, but this also casts in a poor light the entire report.

Instead of a conclusion ..

… Subjective views and wish list on the evolution of the field,
HPC communities, in general, and
within our domains …

….from Tobias

- More methodological high risk-high gain
- Algorithmical novelty in AI proposals
- Open data != reproducible sciences
- Spread out resource requirements evenly
- Parallel submissions to multiple panels → clear differentiation and different focus

….from Maria Paola

Triumph of Artificial Intelligence!

In past years, two main issues:

-proposals in different domains 'spiced up' with AI

-AI - focussed proposals difficult to compare with others

Now:

Dedicated AI track – killing two piegeons – the two issues – with one stone??

- Great rise of astrophysics and cosmology, fuelled by Gravitational Waves
- An increasing interest in plasma science, including magnetic fusion
- Lattice Quantum Chromodynamics – computational nuclear and hadronic physics – no longer dominating the scene but still requesting significant resources

A common computational aspect: the need to explore a large set of parameters, naturally leading to continuation projects, often spanning many years…

Discussion are starting on community access ….
..many pros and cons..

Your views??

(on this and all the other topics we touched upon today!)

THANK YOU

For more information, feel free to visit our website and social media:

/eurohpc-ju.europa.eu /EuroHPC_JU /eurohpc-ju /eurohpc-ju

EuroHPC Joint Undertaking

Generalitat de Catalunya
Departament de Recerca
i Universitats

GOBIERNO DE ESPAÑA
MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES

UPC
UNIVERSITAT POLITÈCNICA DE CATALUNYA
BARCELONATECH

Plan de Recuperación, Transformación y Resiliencia

UNIÓN EUROPEA
Fondo Europeo de Desarrollo Regional

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

RES RED ESPAÑOLA DE SUPERCOMPUTACIÓN

Infraestructuras Científicas y Técnicas Singulares

EuroHPC Joint Undertaking

MareNostrum 5 System and Support

Sergi Girona & David Vicente

Operations Department

sergi.girona@bsc.es
david.vicente@bsc.es

The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Connecting Europe Facility and the Horizon 2020 research and innovation programme, as well as the Participating States Spain, Portugal and Türkiye

EuroHPC Joint Undertaking

Hosting Consortium:

Spain Portugal Türkiye

European Commission

October 2024

MareNostrum 5

Total peak performance: **315.2** Pflops		
General Purpose Partition:	46.4 Pflops	(29-04-2024)
Accelerated Partition:	260 Pflops	(29-04-2024)
Next Generation GPP:	2.82 Pflops	(WiP)
Next Generation ACC:	6 Pflops	(TBA)

MareNostrum 1
2004 – 42.3 Tflops
1st Europe / 4th World
New technologies

MareNostrum 2
2006 – 94.2 Tflops
1st Europe / 5th World
New technologies

MareNostrum 3
2012 – 1.1 Pflops
12th Europe / 36th World

MareNostrum 4
2017 – 11.1 Pflops
2nd Europe / 13th World
New technologies

MareNostrum 5
2022
260 + 46.4 Pflops
8th and 19th World
3rd and 7th Europe

Generalitat de Catalunya
Departament de Recerca i Universitats
Gobierno de España · Ministerio de Ciencia, Innovación y Universidades
UPC Universitat Politècnica de Catalunya BarcelonaTech
Plan de Recuperación, Transformación y Resiliencia
Unión Europea
Fondo Europeo de Desarrollo Regional

MareNostrum5

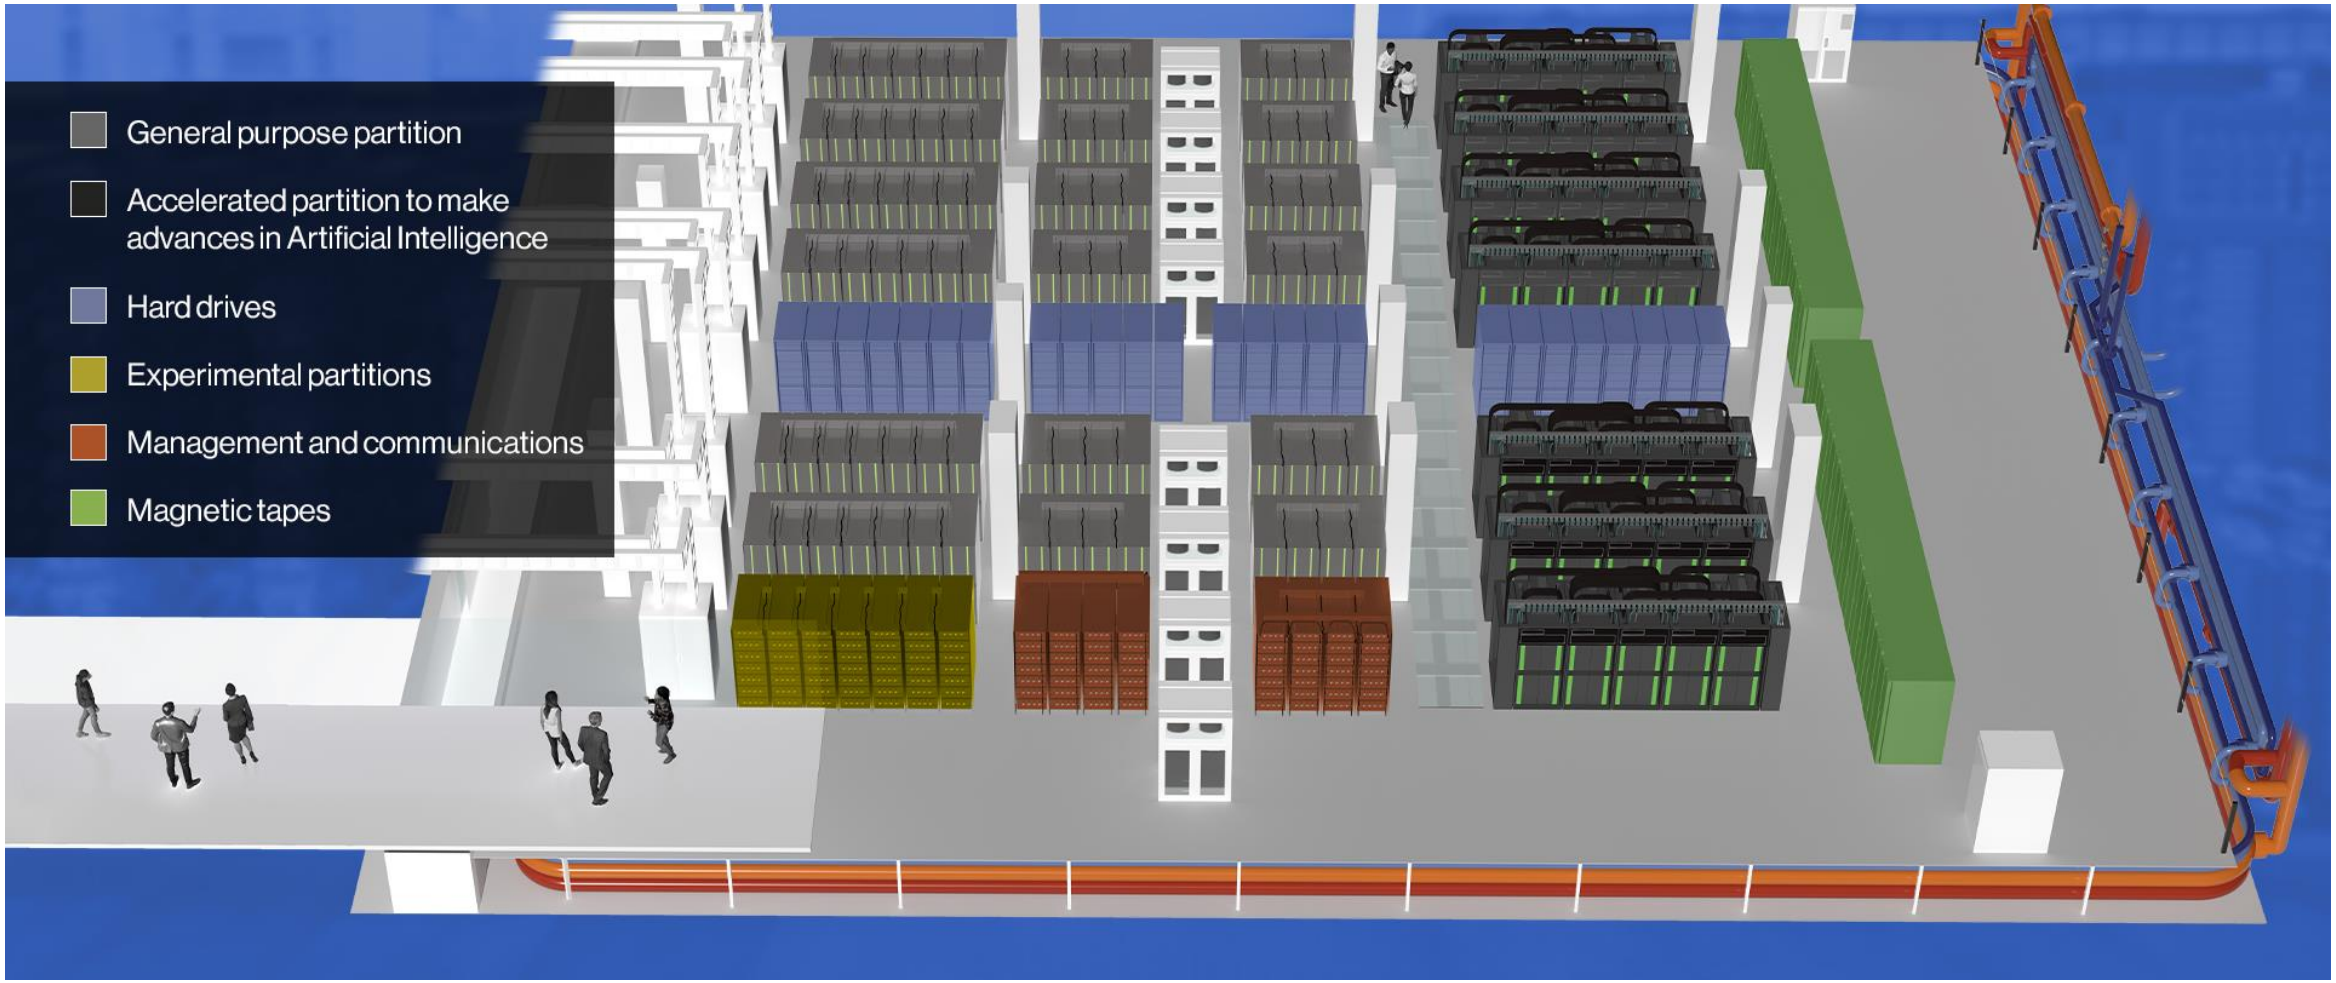

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Top 23, June 2024

Rank	Name	Country	Cores	Accelerators	Rmax [TFlop/s]	Rpeak [TFlop/s]	GFlops/Watts	Computer	Site
1	**Frontier**	**United States**	8.699.904	8.138.240	1.206.000	**1.714.814**	52,93	HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	DOE/SC/Oak Ridge National Laboratory
2	**Aurora**	**United States**	9.264.128	8.159.232	1.012,000	**1.980.006**	26.15	HPE Cray EX - Intel Exascale Compute Blade, Xeon CPU Max 9470 52C 2.4GHz, Intel Data Center GPU Max, Slingshot-11	DOE/SC/Oak Ridge National Laboratory
3	**Eagle**	**United States**	2.073.600	1.900.800	561.200	**846.835**		Microsoft NDv5, Xeon Platinum 8480C 48C 2GHz, NVIDIA H100, NVIDIA Infiniband NDR	Microsoft Azure
4	**Fugaku**	**Japan**	7.630.848		442.010	**537.212**	14,78	Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D	RIKEN Center for Computational Science
5	**LUMI**	**Finland**	2.752.704	2.566.080	379.700	**531.505**	53,43	HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	EuroHPC/CSC
6	**Alps**	**Switzerland**	1.305.600	844.800	270.000	**353.748**	51,98	HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	CSCS
7	**Leonardo**	**Italy**	1.824.768	1.714.176	241.200	**306.311**	32,19	BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, 4x NVIDIA 100	EuroHPC/CINECA
8	**MareNostrum 5 ACC**	**Spain**	663.040	591.360	175.300	**249.435**	42,15	BullSequana XH3000, Xeon Platinum 8460Y+ 40C 2.3GHz, NVIDIA H100 64GB, Infiniband NDR200	EuroHPC/BSC
9	**Summit**	**United States**	2.414.592	2.211.840	148.600	**200.795**	14,72	IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband	DOE/SC/Oak Ridge National Laboratory
10	**Eos NVIDIA DGX SuperPOD**	**United States**	485.888	439.296	121.400	**188.645**		NVIDIA DGX H100, Xeon Platinum 8480C 56C 3.8GHz, NVIDIA H100, Infiniband NDR400	NVIDIA Corporation
11	**Venado**	**United States**	481.440	311.520	98.510	**130.44**	58,29	HPE Cray EX254n, NVIDIA Grace 72C 3.1GHz, NVIDIA GH200 Superchip, Slingshot-11	DOE/NNSA/LANL
12	**Sierra**	**United States**	1.572.480	1.382.400	94.640	**125.712**	12,72	IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband	DOE/NNSA/LLNL
13	**Sunway TaihuLight**	**China**	10.649.600		93.015	**125.436**	6,05	Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway	National Supercomputing Center in Wuxi
14	**Perlmutter**	**United States**	888.832	774.144	79.230	**112.998**	26,90	HPE Cray EX 235n, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-11	DOE/SC/LBNL/NERSC
15	**Selene**	**United States**	555.520	483.840	63.460	**79.215**	23,98	NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband	NVIDIA Corporation
16	**Tianhe-2A**	**China**	4.981.760	4.554.752	61.444	**100.679**	3,32	TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000	National Super Computer Center in Guangzhou
17	**CEA-HE**	**France**	389.232	251.856	57.110	**112.560**	47,32	BullSequana XH3000, Grace Hopper Superchip 72C 3GHz, NVIDIA GH200 Superchip, Quad-Rail BXI v2	Commissariat a l'Energie Atomique (CEA)
18	**Explorer-WUS3**	**United States**	445.440	422.400	53.960	**86.987**		ND96_amsr_MI200_v4, AMD EPYC 7V12 48C 2.45GHz, AMD Instinct MI250X, Infiniband HDR	West US3
19	**ISEG**	**Netherlands**	218.880	200.640	46.540	**86.792**	35,26	Gigabyte G593-SD0, Xeon Platinum 8468 48C 2.1GHz, NVIDIA H100 SXM5 80 GB, Infiniband NDR400	Nebius
20	**Adastra**	**France**	319.072	297.440	46.100	**61.608**	50,03	HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	GENCI-CINES
21	**JUWELS Booster Module**	**Germany**	449.280	404.352	44.120	**70.980**	25,01	Bull Sequana XH2000 , AMD EPYC 7402 24C 2.8GHz, NVIDIA A100, Mellanox HDR InfiniBand/ParTec ParaStation ClusterSuite	Forschungszentrum Juelich (FZJ)
22	**MareNostrum 5 GPP**	**Spain**	725.760		40.102	**46.371**	6,97	ThinkSystem SD650 v3, Xeon Platinum 03H-LC 56C 1.7GHz, Infiniband NDR200	EuroHPC/BSC
23	**Shaheen III - CPU**	**Saudi Arabia**	877.824		35.658	**39.607**	6,73	HPE Cray EX, AMD EPYC 9654 96C 2.4GHz, Slingshot-11	King Abdullah University of Science and Technology

MareNostrum 5 – Storage Spectrum Scale

MN5 Technology providers:
Eviden, Lenovo, IBM, Nvidia, Intel

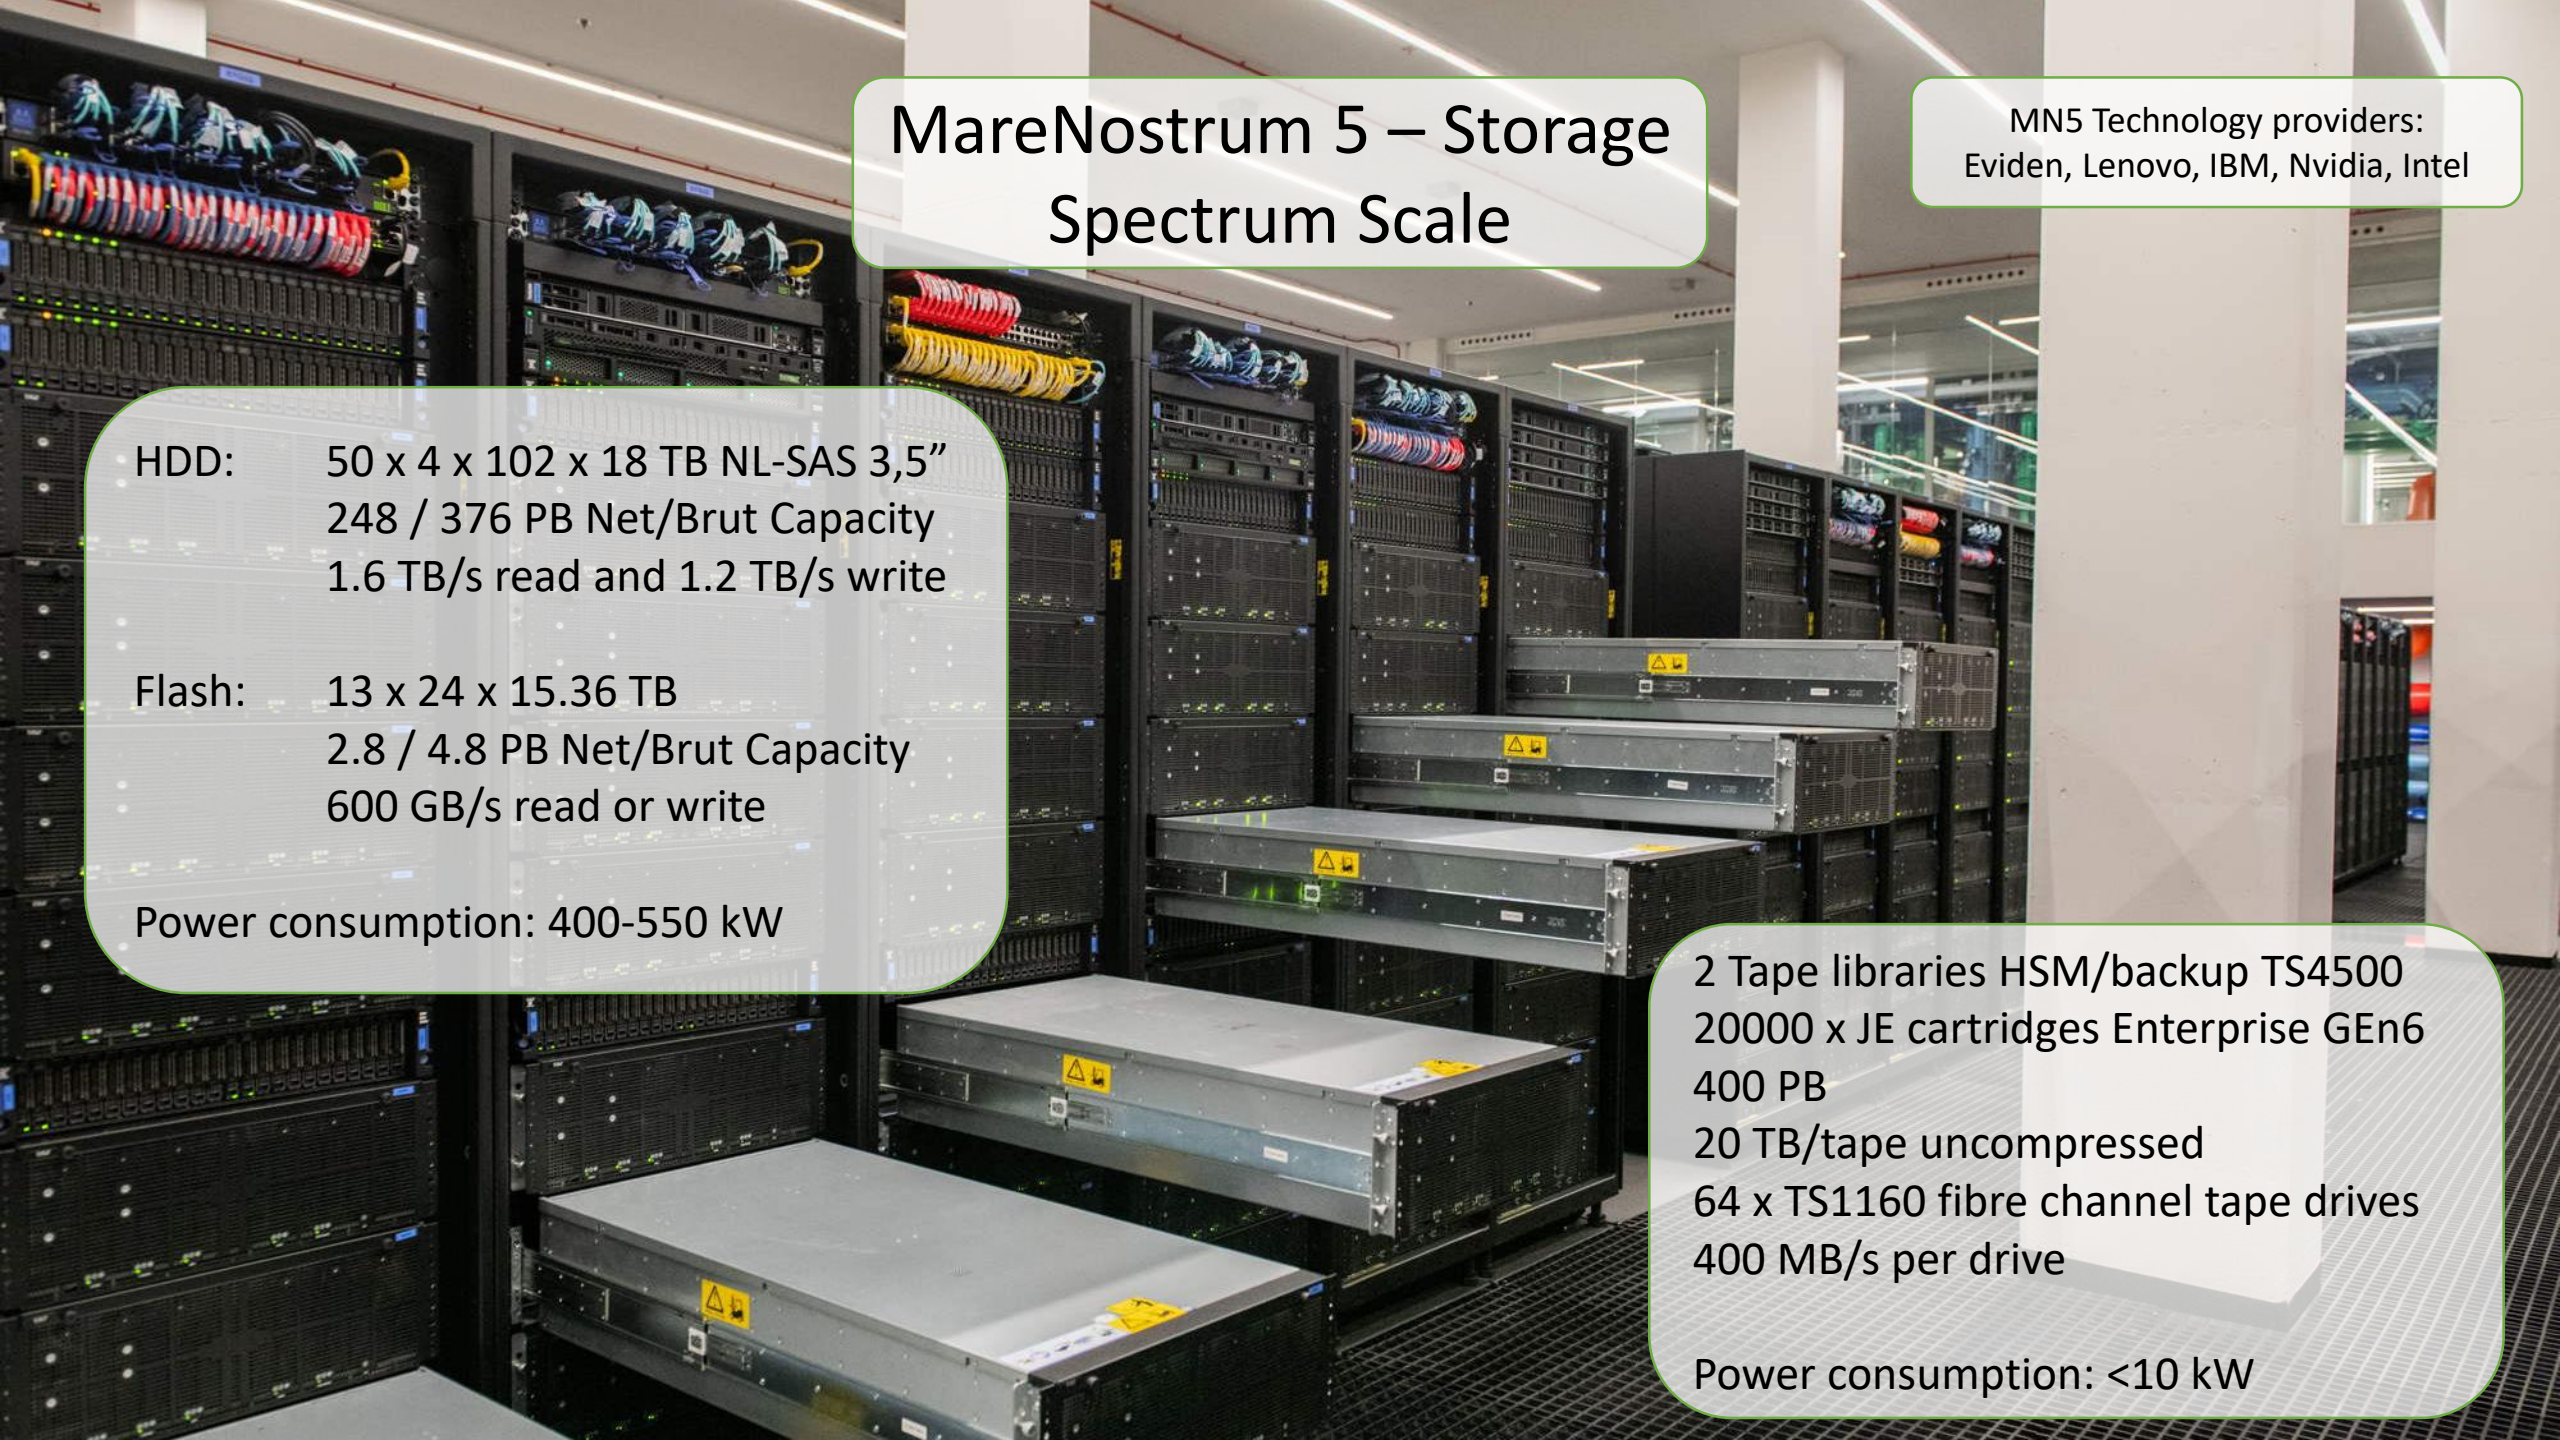

HDD: 50 x 4 x 102 x 18 TB NL-SAS 3,5”
248 / 376 PB Net/Brut Capacity
1.6 TB/s read and 1.2 TB/s write

Flash: 13 x 24 x 15.36 TB
2.8 / 4.8 PB Net/Brut Capacity
600 GB/s read or write

Power consumption: 400-550 kW

2 Tape libraries HSM/backup TS4500
20000 x JE cartridges Enterprise GEn6
400 PB
20 TB/tape uncompressed
64 x TS1160 fibre channel tape drives
400 MB/s per drive

Power consumption: <10 kW

MareNostrum5 GPP

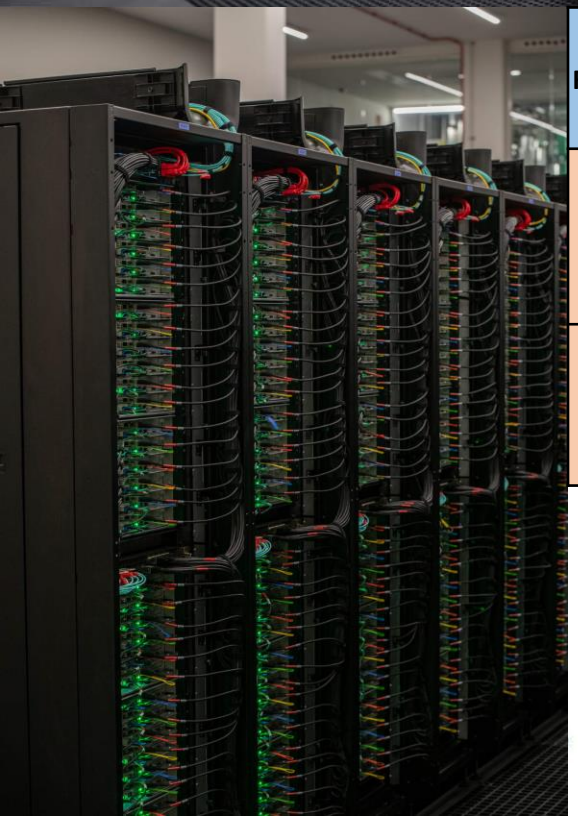

Racks	Cooling	Nodes		Processor/Accelerator		Memory	PFlops (HPL)	Local Drive	High-Perf. Network
		Total	per rack						
89	DLC +RDHX	6192	72 (6x6x2)	2x Intel Sapphire R. 8480+	56c @ 2GHz	>2GB/core 256GB DDR5	40.10	960GB NVMe	1x NDR200 Shared by 2 nodes
		216				>8GB/core 1024GB DDR5			
1		72		2x Intel Sapphire R. 9480	56c @ 1.9GHz	> 0.5GB HBM/core 128GB HBM + 32GB DDR5	0.34		

November 2023

HPL: #19 , #1 x86
HPCG: #24
Green500: #81
5.7 MW under HPL

HPL: 40.10 PFlops
HPCG: 484.36 TFlops
Green500: 6.97 Gflops/watt

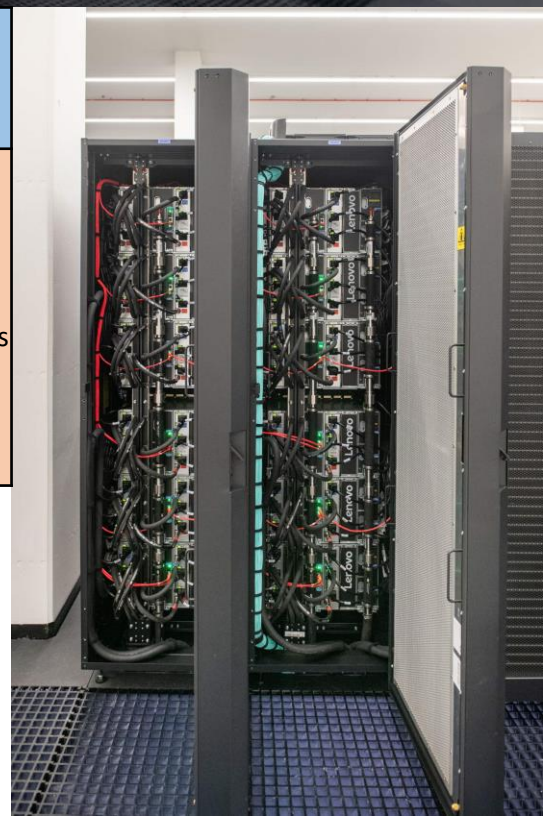

MareNostrum5 ACC

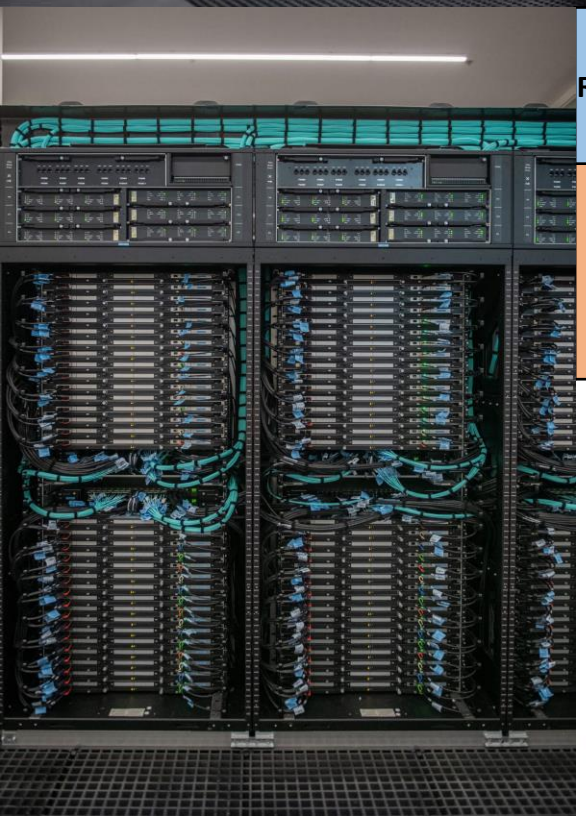

Racks	Cooling	Nodes		Processor/Accelerator		Memory	PFlops (HPL)	Local Drive	High-Perf. Network
		Total	per rack						
35	DLC	1120	32	2x Intel Sapphire R. 8460Y+	40c @ 2GHz	512GB	*175.3*	480GB NVMe	4x NDR200
				4x Nvidia Hopper 64GB HBM					

November 2023
HPL: #8
Green500: #5
2.5 MW under HPL

HPL: 175.30 Pflops
HPCG: 1.146 PFlops
Green500: 42.15 Gflops/watt
Graph500: 15.73e+12 GTEPS (BFS)
4.1 MW under HPL (175 PF/s)

Nvidia Hopper BF16: 120 TFLOPs, Tensorcore: 1 PFLOPs
MN5 Total performance BF16: 536 PFLOPs, Tensorcore: 4.48 EFLOPs

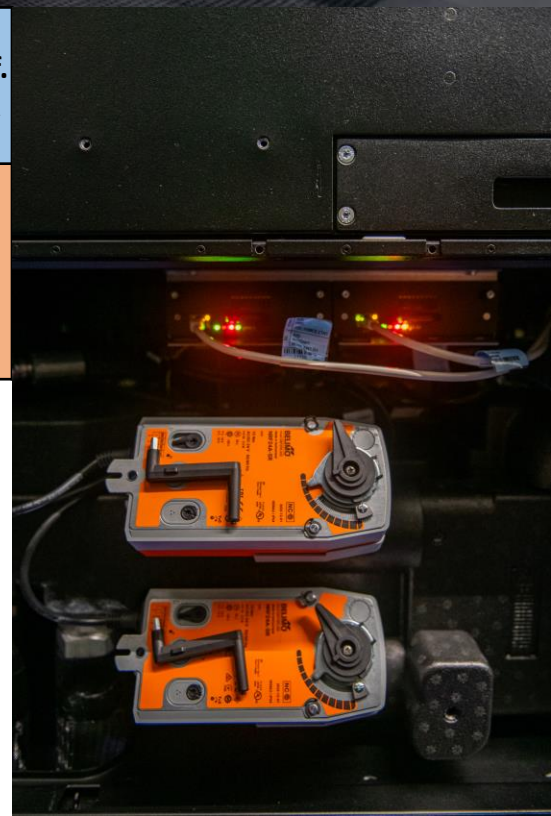

MareNostrum5 NG GP

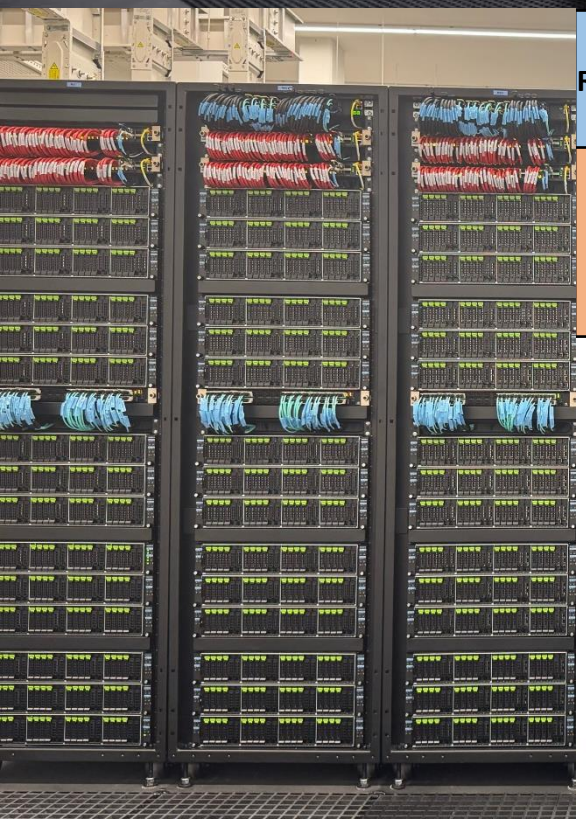

Racks	Cooling	Nodes		Processor/Accelerator		Memory	PFlops (HPL)	Local Drive	High-Perf. Network
		Total	per rack						
7	RDHX	408	6x60 1x48	2x NVidia Grace	72c @ 2GHz	240 GB LPDDR5	*2*	745GB NVMe	1x NDR200

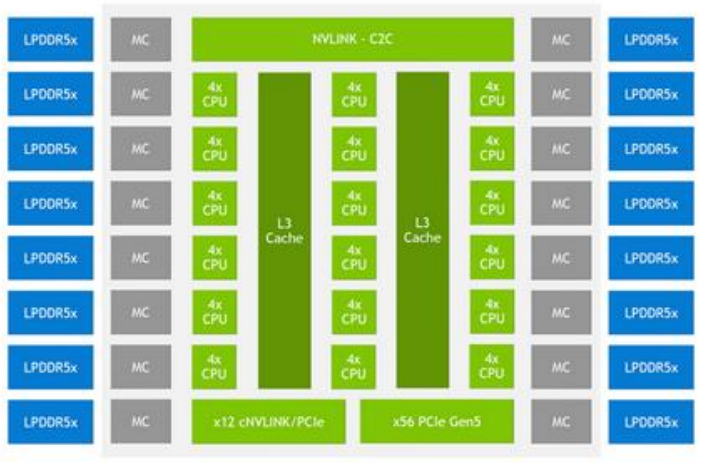

MareNostrum 5 – High Speed network
IB NDR200

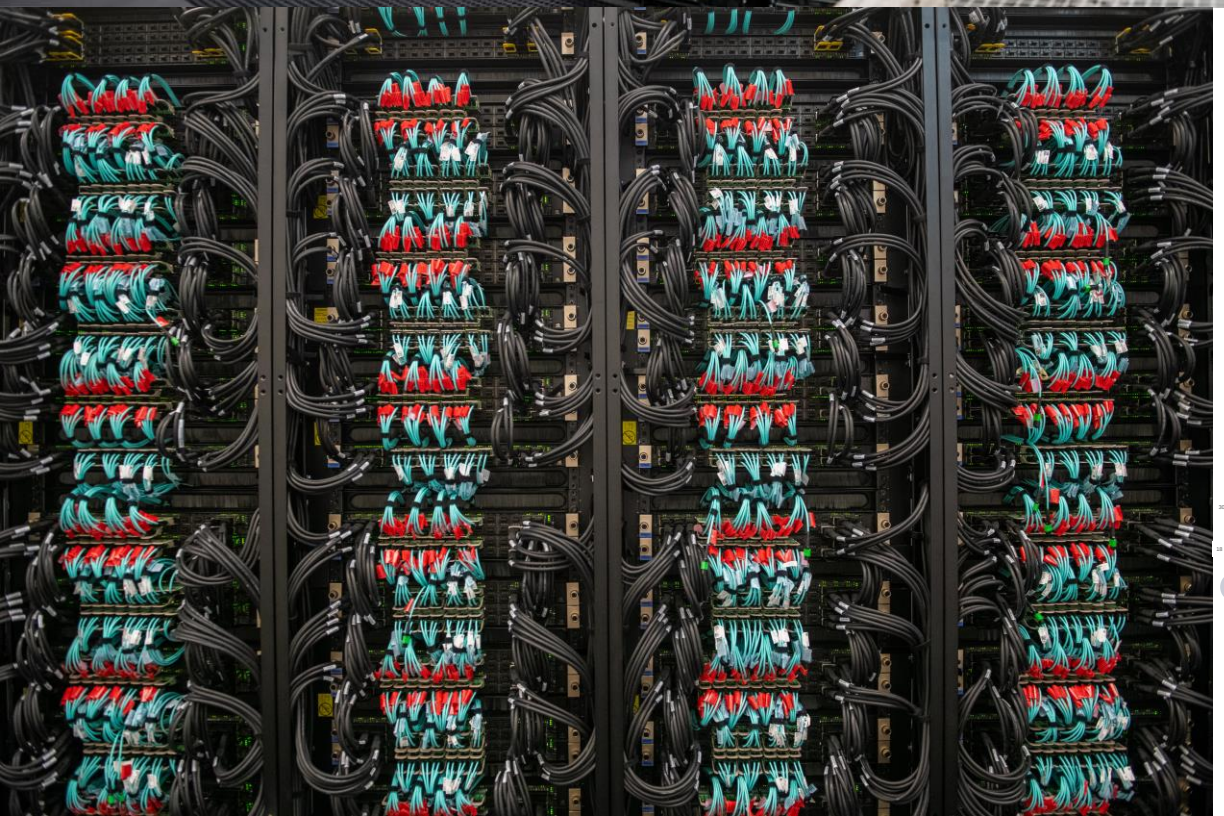

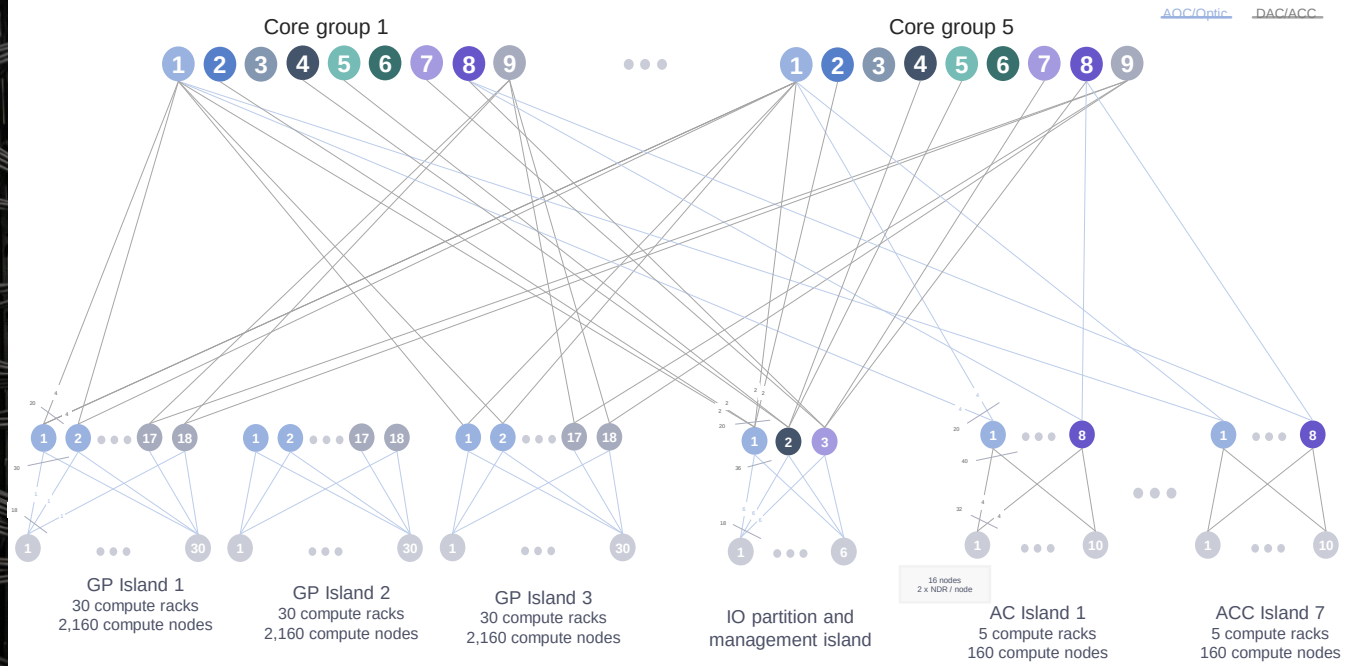

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

RES
RED ESPAÑOLA DE SUPERCOMPUTACIÓN

Infraestructuras Científicas y Técnicas Singulares

EuroHPC Joint Undertaking

MareNostrum 5
Users and Support

The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Connecting Europe Facility and the Horizon 2020 research and innovation programme, as well as the Participating States Spain, Portugal and Türkiye

EuroHPC Joint Undertaking

Hosting Consortium:

Spain Portugal Türkiye

European Commission

MareNostrum5 – specific features

- Ideal for highly scalable applications, being an exceptionally large General-Purpose Processor (GPP) machine with over 6400 nodes and 717,000 cores.

- Well-suited for Large Language Models and Artificial Intelligence, thanks to its over 4400 NVIDIA H100 GPUs (each with 64GB HBM2e memory).

- Excellent for heterogeneous executions, with three partitions (GPP, HBM, and ACC) utilizing the same batch system, allowing mixed jobs.

- Perfect for large-scale data applications, equipped with the IBM Spectrum Scale file system supporting up to 240 petabytes with 1.2 TB/s write and 1.6TB/s read bandwidth and 400 PB on tapes.

MareNostrum 5 Numbers

	EuroHPC	Total
Number of Projects	118	575
Numer of Users	319	1860
Number of jobs submited		> 10 M

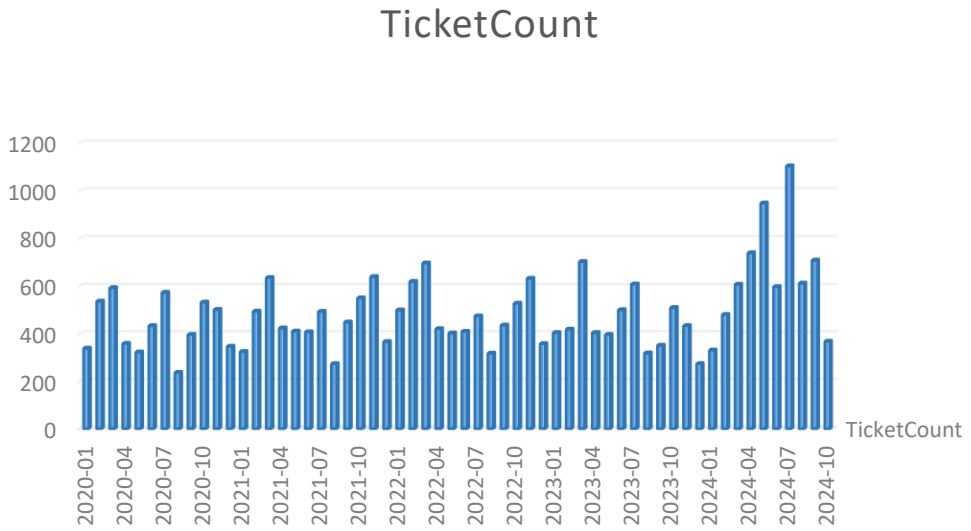

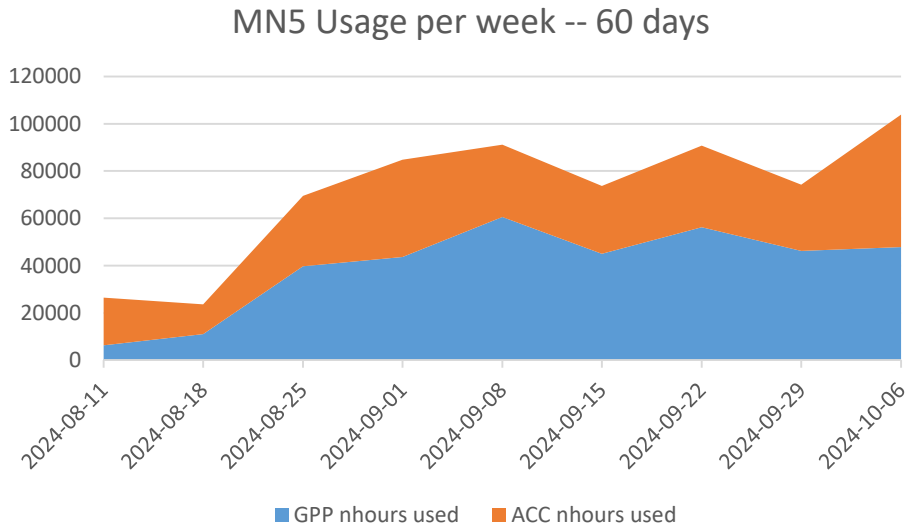

	Installed apps
GPP partition	275
ACC partition	113
TOTAL	388

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

MareNostrum 5 – Important points

Topic of interest	description
Internet access from login nodes	No outgoing connection from any compute or login node This limitation affects the creation of python envs The solution currently implemented is just ask support@bsc.es to install the env for you (as user or in a generic python module)
Containers options available	Only **Singularity** containers are available, and without fakeroot, any NGC container from NVIDIA version 23.7 or higher can be installed and run. Previous versions may cause issues with GLIBC.
LLM models	Good scalability with up to 64 nodes (256 GPUs), for larger runs the filesystem can become a limitation, but 128 nodes has still a quite good efficiency but starts degradation. We are studying ways to improve scalability further.
Energy Efficiency and Power Management	We provide EAR monitoring in all the nodes, ACC and GPP

MareNostrum5 – Important points

Topic of interest	description
Data Transfer nodes for large data movements	200GB/s connection shared between 4 data transfer nodes. SSH protocol and Grid FTP available (GridFTP only under petition)
Network GPP and ACC	GPP – 100GB/s per node ACC -- 800GB/s per node (4 devices of 200GB/s, 1 per GPU) <u>contention:</u> island GPP 2160 nodes (241.920 cores), then 3:2 island ACC 160 nodes (640 GPUs), then 2:1
Userportal	Job status on real time and post-mortem Data from GPUs available and power usage per node

UserPortal

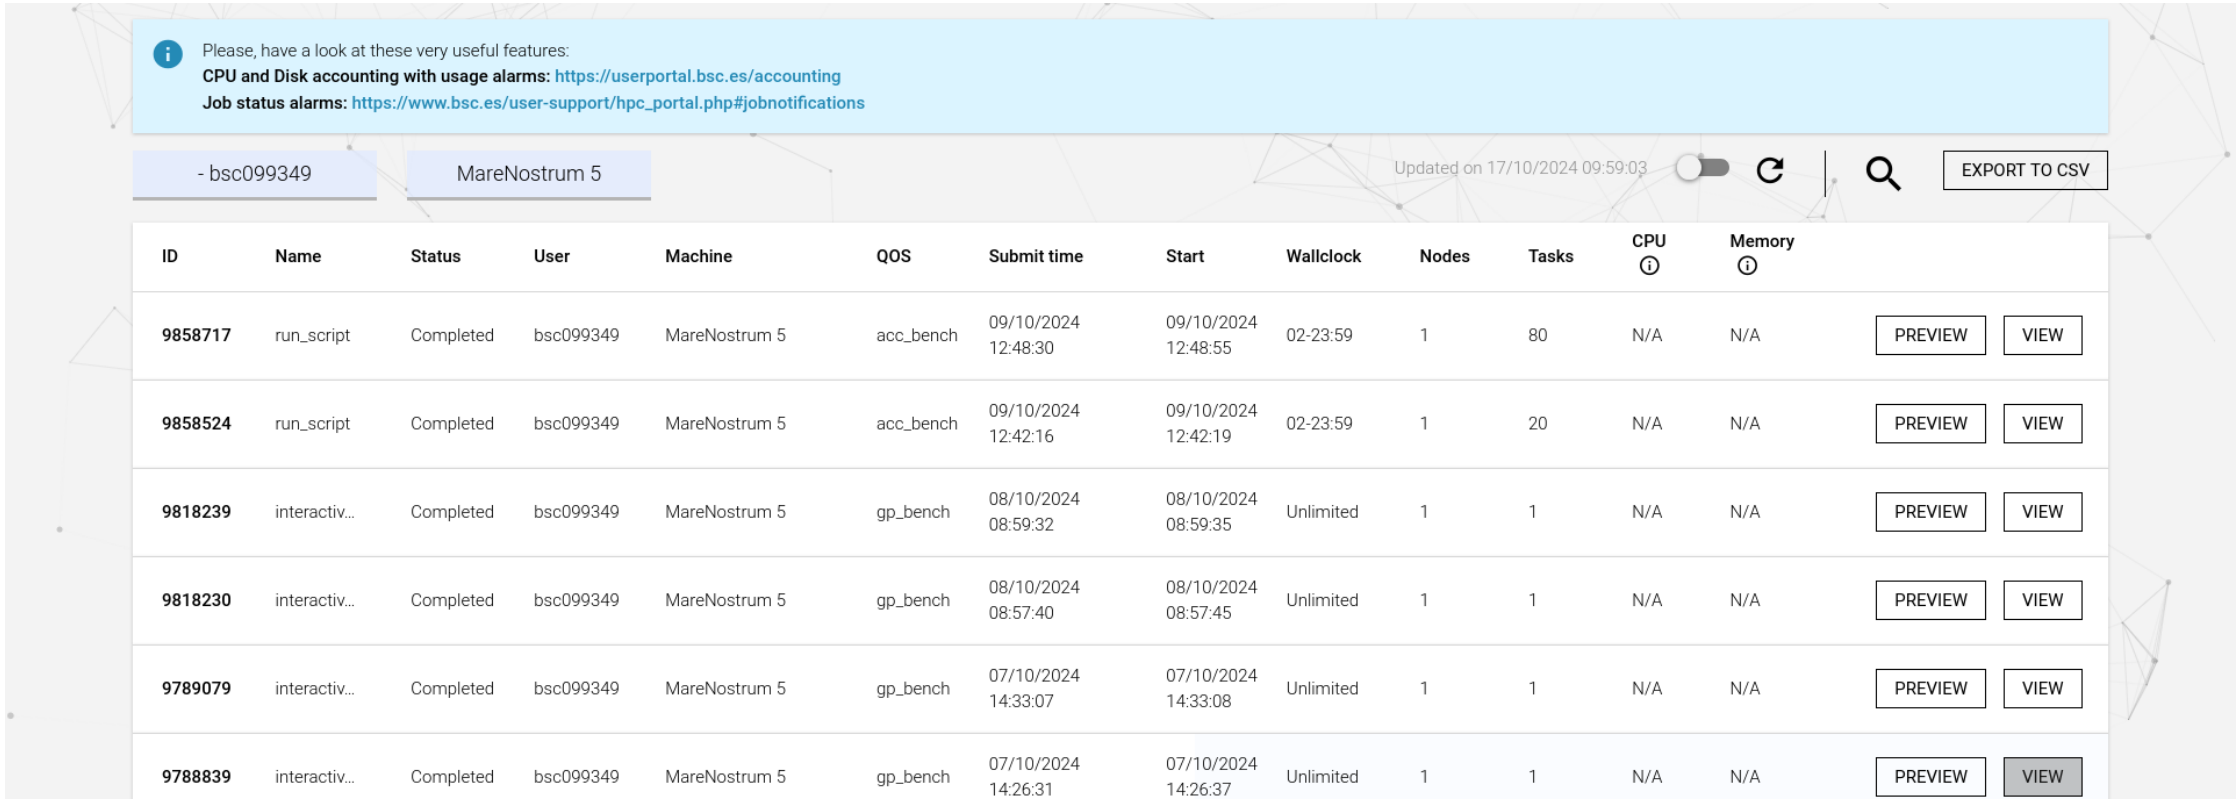

Please, have a look at these very useful features:
CPU and Disk accounting with usage alarms: https://userportal.bsc.es/accounting
Job status alarms: https://www.bsc.es/user-support/hpc_portal.php#jobnotifications

- bsc099349 | MareNostrum 5

Updated on 17/10/2024 09:59:03 | EXPORT TO CSV

ID	Name	Status	User	Machine	QOS	Submit time	Start	Wallclock	Nodes	Tasks	CPU	Memory	
9858717	run_script	Completed	bsc099349	MareNostrum 5	acc_bench	09/10/2024 12:48:30	09/10/2024 12:48:55	02-23:59	1	80	N/A	N/A	PREVIEW VIEW
9858524	run_script	Completed	bsc099349	MareNostrum 5	acc_bench	09/10/2024 12:42:16	09/10/2024 12:42:19	02-23:59	1	20	N/A	N/A	PREVIEW VIEW
9818239	interactiv...	Completed	bsc099349	MareNostrum 5	gp_bench	08/10/2024 08:59:32	08/10/2024 08:59:35	Unlimited	1	1	N/A	N/A	PREVIEW VIEW
9818230	interactiv...	Completed	bsc099349	MareNostrum 5	gp_bench	08/10/2024 08:57:40	08/10/2024 08:57:45	Unlimited	1	1	N/A	N/A	PREVIEW VIEW
9789079	interactiv...	Completed	bsc099349	MareNostrum 5	gp_bench	07/10/2024 14:33:07	07/10/2024 14:33:08	Unlimited	1	1	N/A	N/A	PREVIEW VIEW
9788839	interactiv...	Completed	bsc099349	MareNostrum 5	gp_bench	07/10/2024 14:26:31	07/10/2024 14:26:37	Unlimited	1	1	N/A	N/A	PREVIEW VIEW

UserPortal node compute usage and power consumption

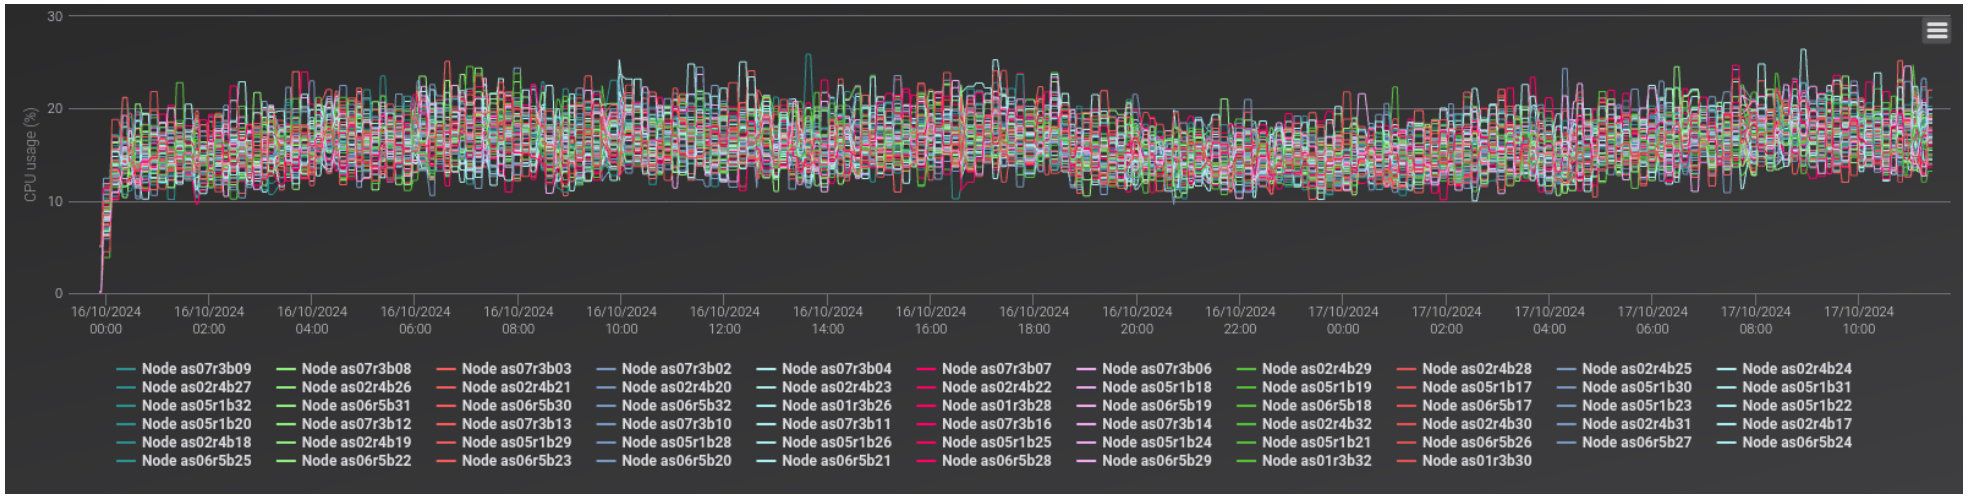

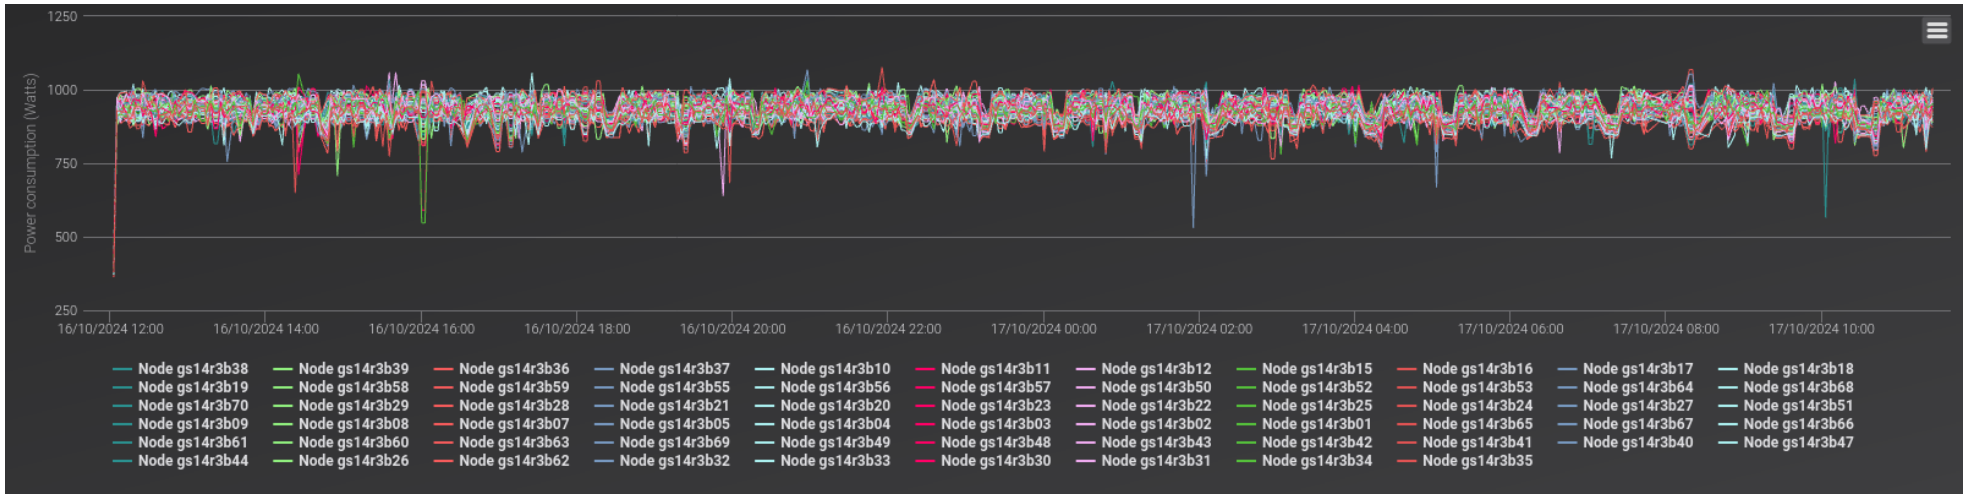

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

UserPortal node compute usage and power consumption

UserPortal GPU0 memory and compute usage

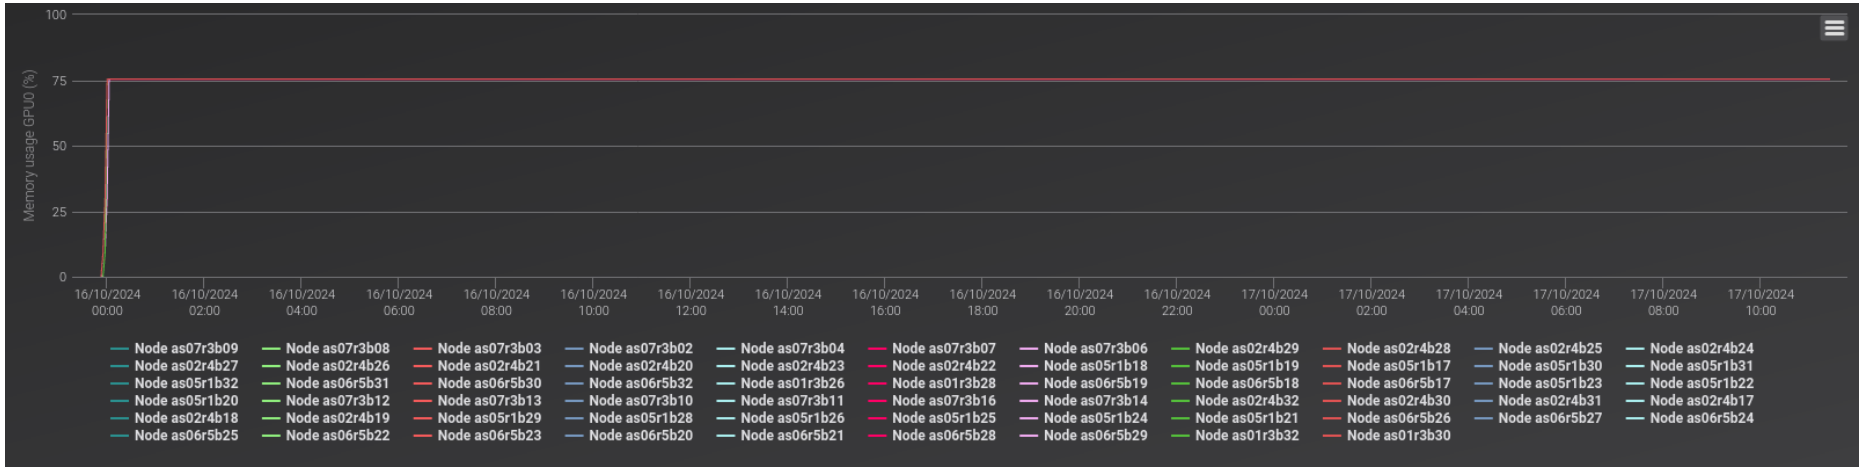

CPU nodes power monitoring

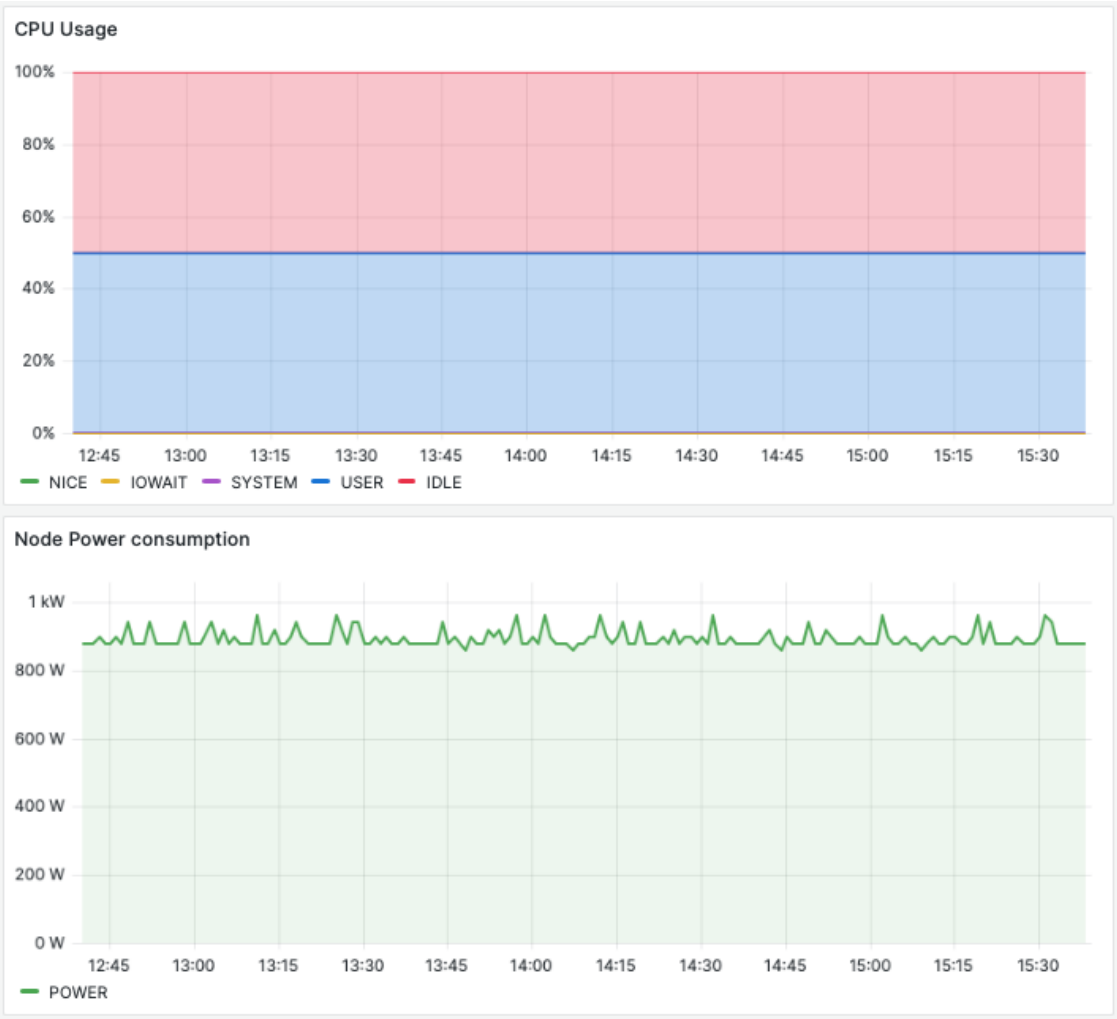

GPU nodes power monitoring

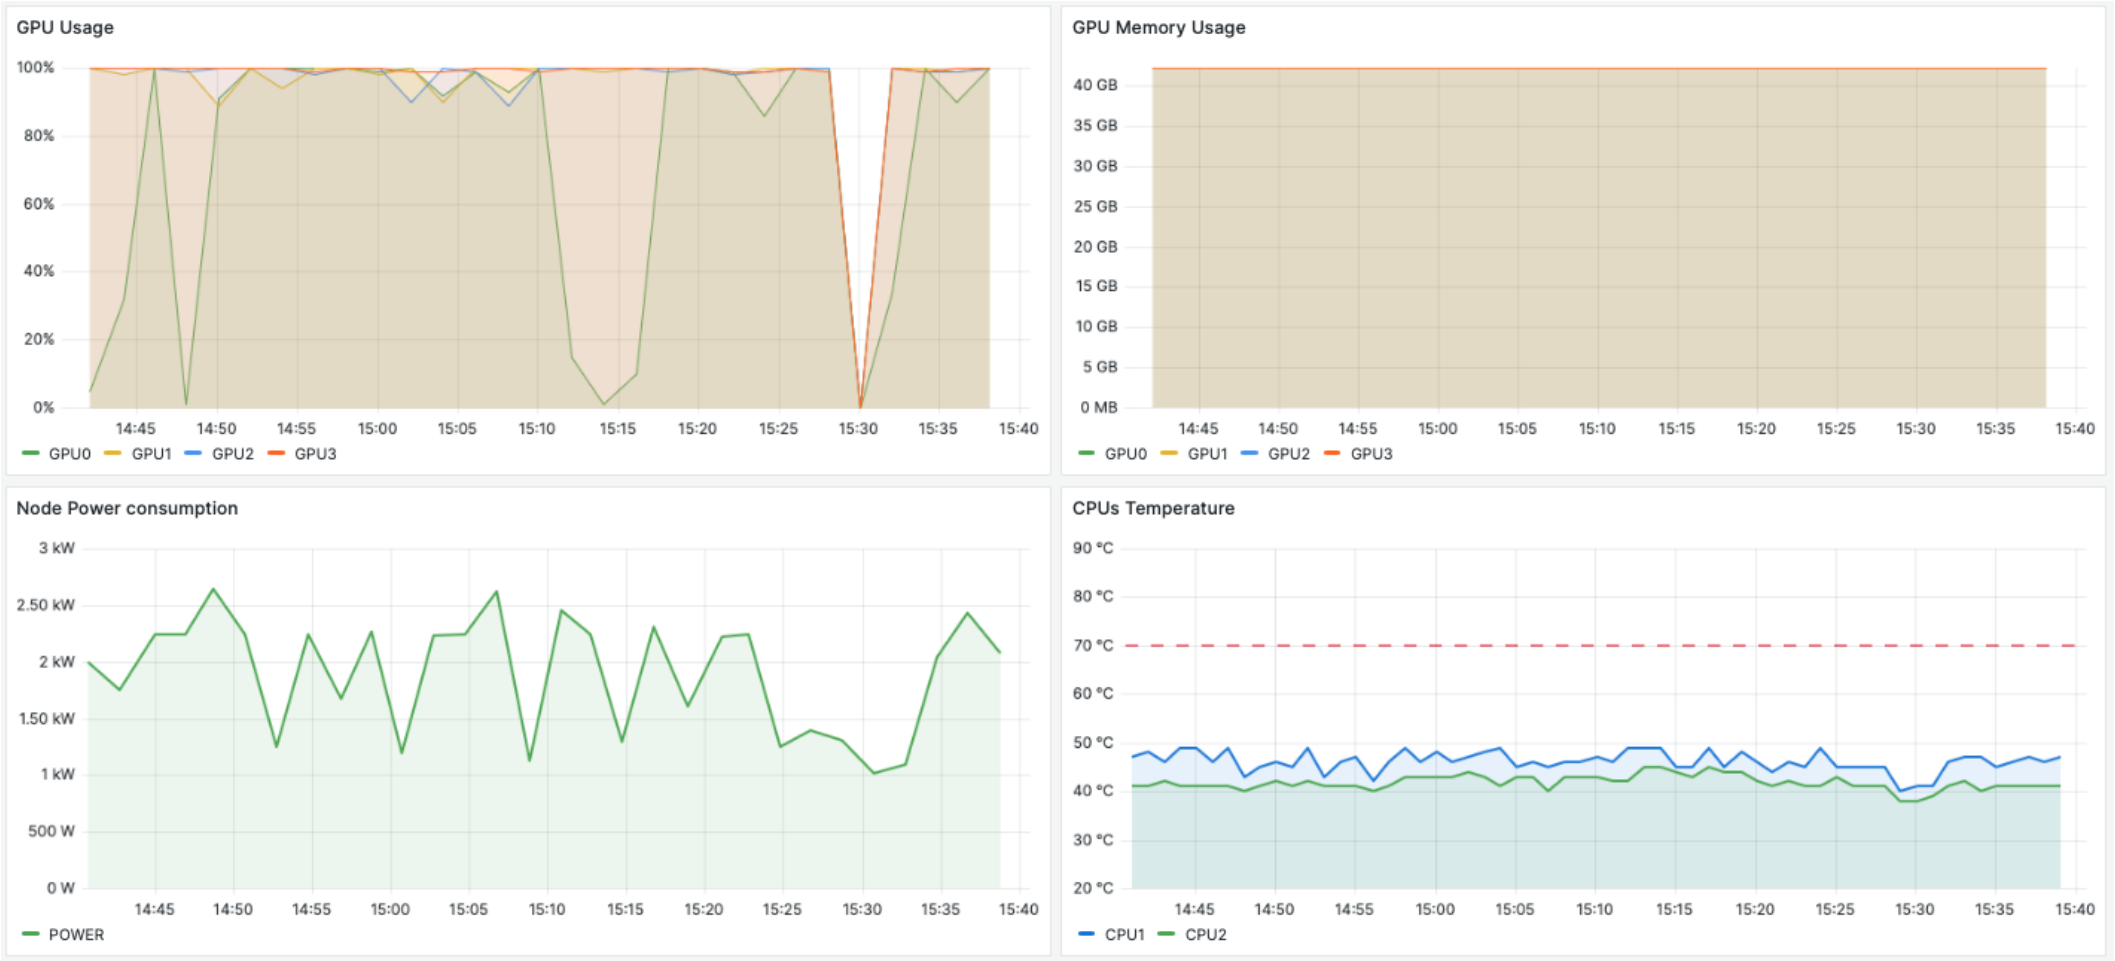

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

MareNostrum5 what is next?

- Quantum partition(s)
  - Quantum Spain and EuroHPC
  - Semiconductors based
- Next Generation Accelerated partition
- AI factories partition
- MareNostrum 6

Thank you

sergi.girona@bsc.es

david.vicente@bsc.es

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

User Behaviour on HPC Vega

Slovenian peta-scale supercomputer

Alja Prah
Jožef Stefan Institute

The acquisition and operation of HPC Vega is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Connecting Europe Facility and the Horizon 2020 research and innovation program, as well as the Participating State Slovenia. The operation HPC RIVR is partly co-funded by the European Union through the European Regional Development Fund and by the Ministry for Science, Education and Sport of the Republic of Slovenia. The operation is carried out within main priority axis no. 1: »International competitiveness of research, innovation and technological development in line with smart specialization for enhanced competitiveness and greening of the economy«, priority investment 1.1 »Enhancing research and innovation (R&I) infrastructure and capacities to develop R&I excellence, and promoting Centres of competence, in particular those of European interest«, specific objective 1.1.1 »Efficient use of the research infrastructure and development of knowledge/competences to improve national and international collaboration in the knowledge triangle« within the Operational Program for the Implementation of the EU Cohesion Policy 2014-2020.

EUROPEAN UNION
COHESION FUND

REPUBLIC OF SLOVENIA
MINISTRY OF EDUCATION, SCIENCE AND SPORT

EuroHPC
Joint Undertaking

System Overview

- 35% EuroHPC JU users + 65% national users
- 2FA (FreeIPA SSH key + OTP)
- long term storage (**Ceph**) + high performance storage (**Lustre**)
- software on our **central CVMFS repo** as EasyBuild modules + EESSI repo
- support for **containers** (--fakeroot, MOFED template)
- own image registry is planned

```
BootStrap: docker
From: quay.io/centos/centos:stream8


%files
/home/barbarak/Downloads/MLNX_OFED_LINUX-5.8-1.1.2.1-rhel8.7-x86_64.tgz /opt/MLNX_OFED_LINUX-5.8-1.1.2.1-rhel8.7-
x86_64.tgz
#/home/barbarak/Downloads/RPM-GPG-KEY-Rocky-8 /etc/pki/rpm-gpg/RPM-GPG-KEY-Rocky-8
#/home/barbarak/Downloads/RPM-GPG-KEY-rockyofficial /etc/pki/rpm-gpg/RPM-GPG-KEY-rockyofficial

%environment
    export OMPI_DIR=/usr/local
    export SINGULARITY_OMPI_DIR=$OMPI_DIR
    export SINGULARITYENV_APPEND_PATH=$OMPI_DIR/bin
    export SINGULAIRTYENV_APPEND_LD_LIBRARY_PATH=$OMPI_DIR/lib

%post

    ## Prerequisites
    dnf -y update
    dnf install --nogpgcheck -y dnf-plugins-core
    dnf config-manager --set-enabled powertools
    dnf groupinstall -y 'Development Tools'
    dnf install --nogpgcheck -y wget git bash hostname gcc gcc-gfortran gcc-c++ make file autoconf automake
libtool zlib-devel python3
    dnf install --nogpgcheck -y libmnl lsof numactl-libs ethtool tcl tk

    ## Packages required for OpenMPI and PMIx
    dnf install --nogpgcheck -y libnl3 libnl3-devel
    dnf install --nogpgcheck -y libevent libevent-devel
    dnf install --nogpgcheck -y munge munge-devel

    # Mellanox OFED matching VEGA
    mkdir -p /opt/mofed
    cd /opt/mofed
    #wget -c https://content.mellanox.com/ofed/MLNX_OFED-5.4-3.1.0.0/MLNX_OFED_LINUX-5.4-3.1.0.0-rhel8.5-
x86_64.tgz
    tar xf ../MLNX_OFED_LINUX-5.8-1.1.2.1-rhel8.7-x86_64.tgz
    cd MLNX_OFED_LINUX-5.8-1.1.2.1-rhel8.7-x86_64
    ./mlnxofedinstall --basic --user-space-only --without-fw-update --distro rhel8.7 --force

    # PMIx
    mkdir -p /tmp/pmix
    cd /tmp/pmix
    wget -c https://github.com/openpmix/openpmix/releases/download/v3.1.5/pmix-3.1.5.tar.gz
    tar xf pmix-3.1.5.tar.gz
    cd pmix-3.1.5
    ./configure --prefix=/usr/local --with-munge=/usr && \
    make -j
    make install

  # libfabric
  mkdir -p /tmp/libfabric
  cd /tmp/libfabric
  wget https://github.com/ofiwg/libfabric/releases/download/v1.17.0/libfabric-1.17.0.tar.bz2
  tar xf libfabric*tar.bz2
  cd libfabric-1.17.0
  ./configure --prefix=/usr/local && \
  make -j
  make install
  ...
```

Common User Behaviour Challenges

VEGA

Most tickets are first level support:

- two-factor authentication (**2FA**) and **SSH** key setup
- understanding resource **priority**

```
Job_priority =
    (PriorityWeightFairshare) * (fair-share_factor) +
    (PriorityWeightAge) * (age_factor) +
    (PriorityWeightJobSize) * (job_size_factor) +
    (PriorityWeightPartition) * (partition_factor) +
    (PriorityWeightQOS) * (QOS_factor) +
    SUM(TRES_weight_cpu * TRES_factor_cpu,
        TRES_weight_<type> * TRES_factor_<type>
```

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

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

EuroHPC Joint Undertaking

REPUBLIC OF SLOVENIA
MINISTRY OF EDUCATION, SCIENCE AND SPORT

HPC RIVR

EUROPEAN UNION
COHESION FUND

Quota & Accounts

- **quota** on home directory and usage of storage
- **multiple projects** usage

Quotas

Quotas	Capacity	Description
home	100GB	Size of home directory for each user
scratch	20GB	Size of scratched directory for each user

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

Multiple Project Usage

Users with multiple project access have multiple Slurm accounts. For each project, the Slurm account is unique. The users on the same project have the same Slurm account.

The default account is usually set on the first approved project and can be changed by sending a request to `support@sling.si`.

Users can get information of their Slurm accounts with commands:

```
sacctmgr show assoc cluster=vega user=$USER format=account%40 -n
```

This is very important, because users with multiple project access in order to send jobs on the right project must use the `flag --account` in the `srun` and `sbatch` command.

Using `flag --account` enabling usage of resources on the right project. If the user does not set `flag --account`, then the job will be sent on the default set account, project.

SBATCH example for multiple project access:

```
#!/bin/bash
#SBATCH --job-name=my_job
#SBATCH --account=example
#SBATCH --partition=cpu
#SBATCH --cpus-per-task=12
#SBATCH --mem=32GB
#SBATCH --time=01:00:00
```

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

REPUBLIC OF SLOVENIA
MINISTRY OF EDUCATION,
SCIENCE AND SPORT

EUROPEAN UNION
COHESION FUND

Underutilized Resources & Non-Active Users

Projects with **less than 10%** of allocated resouces used

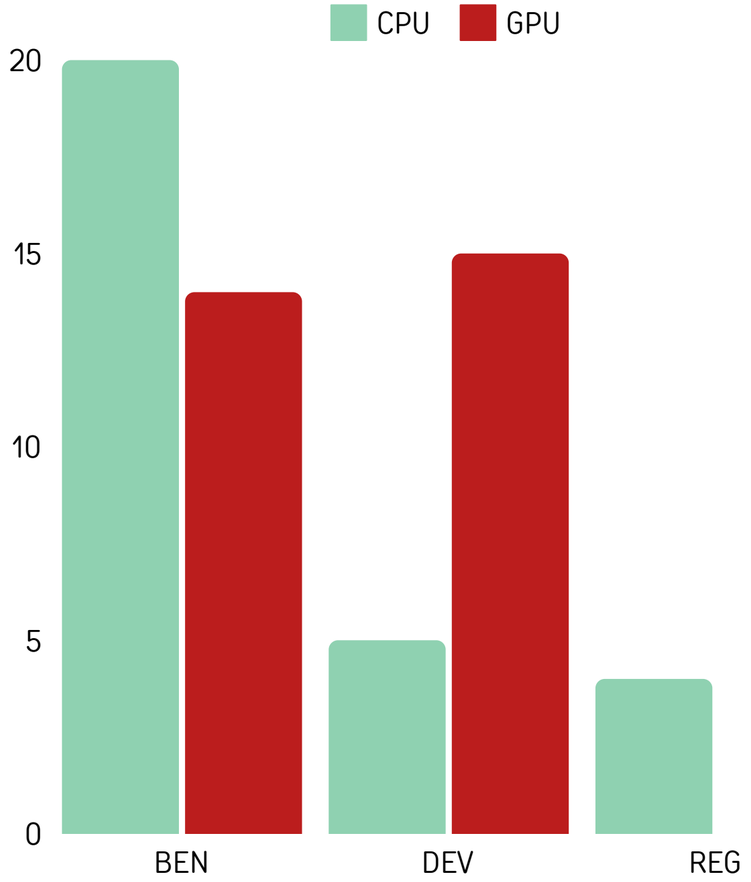

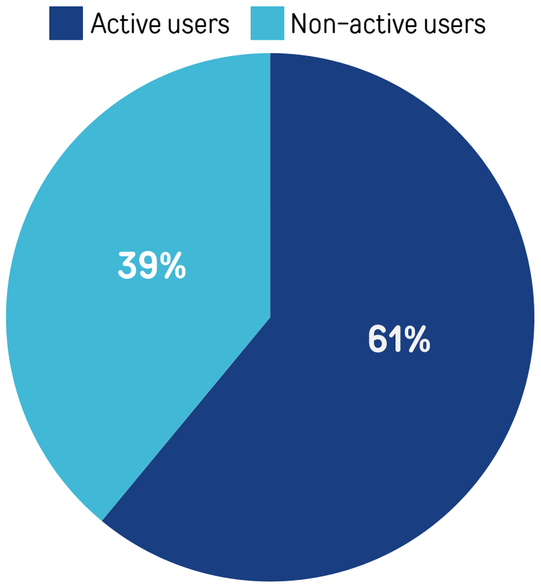

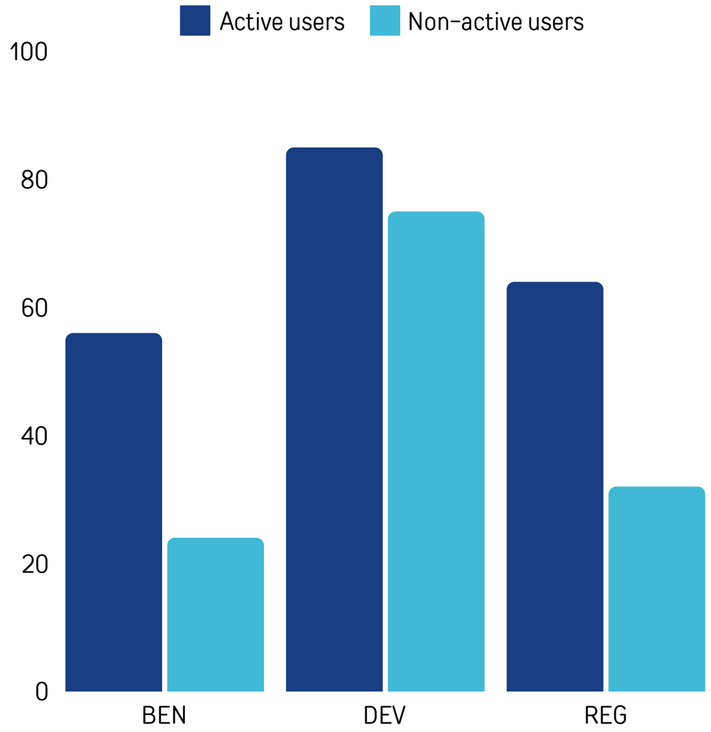

Extensions

AI, REGULAR & EXTREME SCALE ACCESS CALLS:

- **3 months** – no additional resources
- max 20% of the initial allocation
- request at the latest **1 month** before original end of allocation
- send request to **access@eurohpc-ju.europa.eu**

BENCHMARK & DEVELOPMENT ACCESS CALLS:

- extensions **not granted**
- **exceptional** in case of **technical issues** coming from the centre

VEGA

EuroHPC Joint Undertaking

REPUBLIC OF SLOVENIA
MINISTRY OF EDUCATION, SCIENCE AND SPORT

HPC RIVR

EUROPEAN UNION
COHESION FUND

Storage Issues

VEGA

- high performance storage Lustre not suitable for all workflows
- use RAM/local disk/Ceph for scratch
- environments with lots of **small files** (i.e. Conda/pip) = containers recommended

CONDA

ceph

lustre

EuroHPC Joint Undertaking

REPUBLIC OF SLOVENIA
MINISTRY OF EDUCATION, SCIENCE AND SPORT

HPC RIVR

EUROPEAN UNION
COHESION FUND

Software & Submit

- choosing the **correct version** of different **software** (ask!)
- providing **modules**, **job submission** help (MPI, UCX, mem...), container templates

```
FlexiBLAS/3.2.0-GCC-11.3.0                    ParaView/5.11.2-foss-2023a                (D)   gettext/0.21.1-GCCcore-12.3.0                 nsync/1.24.0-GCCcore-11.2.0
FlexiBLAS/3.2.1-GCC-12.2.0                    Perl-bundle-CPAN/5.36.1-GCCcore-12.3.0      gettext/0.21.1                              nsync/1.25.0-GCCcore-11.3.0               (D)
FlexiBLAS/3.3.1-GCC-12.3.0             (D)    Perl/5.28.0-GCCcore-7.3.0                   gettext/0.22                          (D)   numactl/2.0.11-GCCcore-5.4.0
FreeImage/3.18.0-GCCcore-10.3.0               Perl/5.28.1-GCCcore-8.2.0                   gfbf/2022b                                  numactl/2.0.11-GCCcore-6.4.0
FriBidi/1.0.9-GCCcore-9.3.0                   Perl/5.30.0-GCCcore-8.3.0                   gfbf/2023a                            (D)   numactl/2.0.11-GCCcore-7.3.0
FriBidi/1.0.10-GCCcore-10.2.0                 Perl/5.30.2-GCCcore-9.3.0                   giflib/5.2.1-GCCcore-8.3.0                  numactl/2.0.12-GCCcore-8.2.0
FriBidi/1.0.10-GCCcore-10.3.0                 Perl/5.32.0-GCCcore-10.2.0                  giflib/5.2.1-GCCcore-10.2.0                 numactl/2.0.12-GCCcore-8.3.0
FriBidi/1.0.10-GCCcore-11.2.0                 Perl/5.32.1-GCCcore-10.3.0-minimal          giflib/5.2.1-GCCcore-10.3.0                 numactl/2.0.13-GCCcore-9.3.0
FriBidi/1.0.12-GCCcore-11.3.0                 Perl/5.32.1-GCCcore-10.3.0                  giflib/5.2.1-GCCcore-11.2.0                 numactl/2.0.13-GCCcore-10.2.0
FriBidi/1.0.12-GCCcore-12.2.0                 Perl/5.34.0-GCCcore-11.2.0                  giflib/5.2.1-GCCcore-11.3.0                 numactl/2.0.14-GCCcore-10.3.0
FriBidi/1.0.12-GCCcore-12.3.0          (D)    Perl/5.34.1-GCCcore-11.3.0-minimal          giflib/5.2.1-GCCcore-12.2.0           (D)   numactl/2.0.14-GCCcore-11.2.0
GCC/4.8.5                                     Perl/5.34.1-GCCcore-11.3.0                  git/2.23.0-GCCcore-8.3.0-nodocs             numactl/2.0.14-GCCcore-11.3.0
GCC/5.4.0-2.26                                Perl/5.36.0-GCCcore-12.2.0                  git/2.23.0-GCCcore-8.3.0                    numactl/2.0.16-GCCcore-12.2.0
GCC/6.4.0-2.28                                Perl/5.36.1-GCCcore-12.3.0                  git/2.28.0-GCCcore-10.2.0-nodocs            numactl/2.0.16-GCCcore-12.3.0
GCC/7.3.0-2.30                                Perl/5.38.0-GCCcore-13.2.0            (D)   git/2.32.0-GCCcore-10.3.0-nodocs      (D)   numactl/2.0.16-GCCcore-13.2.0             (D)
GCC/8.2.0-2.31.1                              Pillow-SIMD/8.2.0-GCCcore-10.3.0            git/2.33.1-GCCcore-11.2.0-nodocs            occt/7.5.0-foss-2021a
GCC/8.3.0                                     Pillow-SIMD/9.2.0-GCCcore-11.3.0      (D)   git/2.36.0-GCCcore-11.3.0-nodocs      (D)   patchelf/0.18.0-GCCcore-12.3.0
GCC/9.3.0                                     Pillow/8.0.1-GCCcore-10.2.0                 git/2.38.1-GCCcore-12.2.0-nodocs            pixman/0.34.0-GCCcore-6.4.0
GCC/10.2.0                                    Pillow/8.2.0-GCCcore-10.3.0                 git/2.41.0-GCCcore-12.3.0-nodocs      (D)   pixman/0.38.4-GCCcore-9.3.0
GCC/10.3.0                                    Pillow/8.3.2-GCCcore-11.2.0                 glibc/2.29-GCCcore-8.3.0                    pixman/0.40.0-GCCcore-10.2.0
GCC/11.2.0                                    Pillow/9.1.1-GCCcore-11.3.0                 glibc/2.30-GCCcore-8.3.0                    pixman/0.40.0-GCCcore-10.3.0
GCC/11.3.0                                    Pillow/10.0.0-GCCcore-12.3.0          (D)   glibc/2.31-GCCcore-9.3.0              (D)   pixman/0.40.0-GCCcore-11.2.0
GCC/12.2.0                                    PnetCDF/1.12.2-gompi-2020b                  gmpy2/2.1.2-GCC-11.3.0                      pixman/0.40.0-GCCcore-11.3.0
GCC/12.3.0                                    PnetCDF/1.12.2-gompic-2020b                 gmsh/4.10.2-foss-2021a                      pixman/0.42.2-GCCcore-12.2.0
GCC/13.2.0                             (D)    PnetCDF/1.12.3-gompi-2021b                  gnuplot/5.2.8-GCCcore-9.3.0                 pixman/0.42.2-GCCcore-12.3.0              (D)
GCCcore/5.4.0                                 PnetCDF/1.12.3-iimpi-2022a            (D)   gnuplot/5.4.2-GCCcore-10.3.0          (D)   pkg-config/0.29.2-GCCcore-6.4.0
GCCcore/6.2.0                                 PyTorch/1.7.1-fosscuda-2020b                gnuplot/5.4.2-GCCcore-11.2.0                pkg-config/0.29.2-GCCcore-7.3.0
GCCcore/6.4.0                                 PyTorch/1.8.1-foss-2020b                    gnuplot/5.4.4-GCCcore-11.3.0                pkg-config/0.29.2-GCCcore-8.2.0
GCCcore/7.3.0                                 PyTorch/1.9.0-fosscuda-2020b                gnuplot/5.4.8-GCCcore-12.3.0          (D)   pkg-config/0.29.2-GCCcore-8.3.0
GCCcore/8.2.0                                 PyTorch/1.10.0-foss-2021a-CUDA-11.3.1       gobff/2021a                                 pkg-config/0.29.2-GCCcore-9.3.0
GCCcore/8.3.0                                 PyTorch/1.12.0-foss-2022a-CUDA-11.7.0       gompi/2017b                                 pkg-config/0.29.2-GCCcore-10.2.0
GCCcore/9.3.0                                 PyTorch/1.12.1-foss-2022a-CUDA-11.7.0       gompi/2018a                                 pkg-config/0.29.2-GCCcore-10.3.0
GCCcore/10.2.0                                PyTorch/1.13.1-foss-2022a                   gompi/2018b                                 pkg-config/0.29.2-GCCcore-11.2.0          (D)
GCCcore/10.3.0                                PyTorch/2.0.1-foss-2022a              (D)   gompi/2019a                                 pkgconf/1.8.0-GCCcore-11.2.0
GCCcore/11.2.0                                PyYAML/5.3.1-GCCcore-10.2.0                 gompi/2019b                                 pkgconf/1.8.0-GCCcore-11.3.0
GCCcore/11.3.0                                PyYAML/5.4.1-GCCcore-10.3.0                 gompi/2020a                                 pkgconf/1.8.0
GCCcore/12.2.0                                PyYAML/5.4.1-GCCcore-11.2.0                 gompi/2020b                                 pkgconf/1.9.3-GCCcore-10.3.0
```

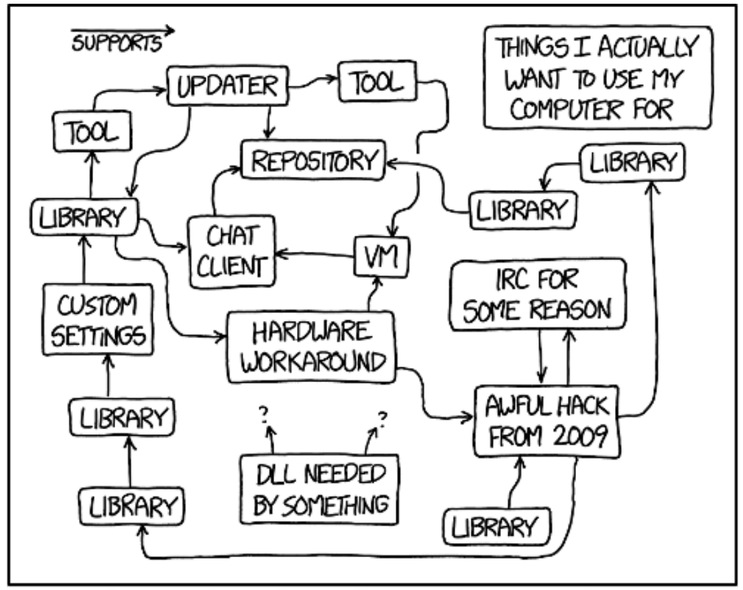

EVERY NOW AND THEN I REALIZE I'M MAINTAINING A HUGE CHAIN OF TECHNOLOGY SOLELY TO SUPPORT ITSELF.

source: https://xkcd.com/1579/

EuroHPC Joint Undertaking

REPUBLIC OF SLOVENIA
MINISTRY OF EDUCATION, SCIENCE AND SPORT

HPC RIVR

EUROPEAN UNION
COHESION FUND

Security

- HPC systems don’t emphasize security enough = too much trust placed in users!
- **zero trust principle** needs to be implemented (malicious internal user poses biggest risk)
- security very important, especially for industry users
- 2FA implemented successfully
- custom kernel – **immediate response** to critical vulnerabilities
- security experts on Vega team
- on-the-fly **risk assesment**

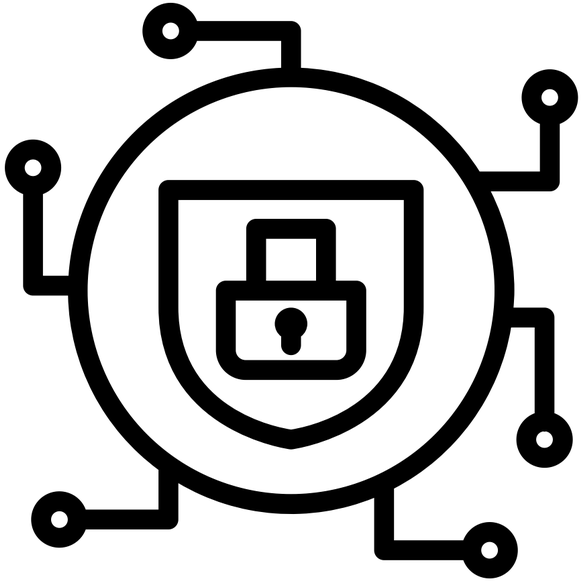

EuroHPC Joint Undertaking
REPUBLIC OF SLOVENIA MINISTRY OF EDUCATION, SCIENCE AND SPORT
HPC RIVR
EUROPEAN UNION COHESION FUND

Projects on HPC Vega

Documentation & Feedback

VEGA

- **documentation**
- sending final report and **feedback** to support team

support@sling.si

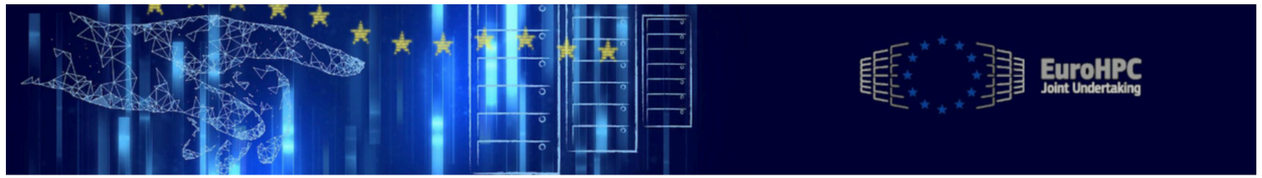

EuroHPC Joint Undertaking (JU) Regular Access

Final Report

REPUBLIC OF SLOVENIA
MINISTRY OF EDUCATION,
SCIENCE AND SPORT

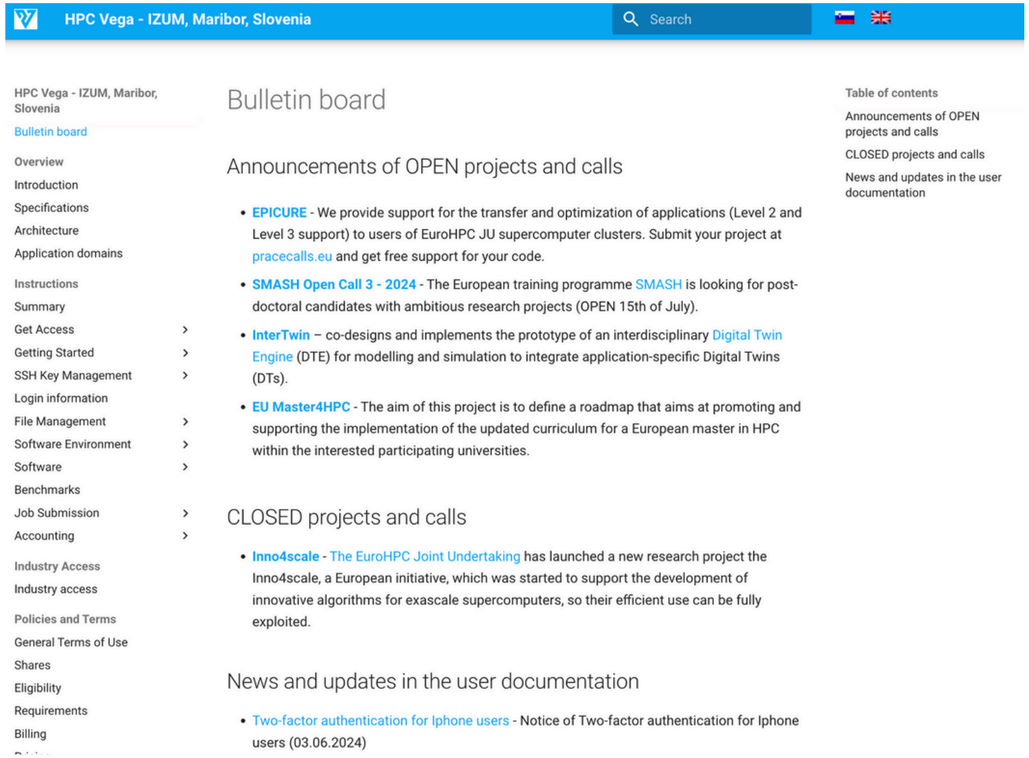

https://en-vegadocs.vega.izum.si/

HPC RIVR

EUROPEAN UNION
COHESION FUND

360°

VEGA VIRTUAL TOUR

RIVR1 – Today

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

SUPPORT CENTRE FOR HPC-POWERED ARTIFICIAL INTELLIGENCE (AI) APPLICATIONS (AISC)

EHPC USER DAY– AMSTERDAM, NL, OCTOBER 22, 2024

MLADEN.SKELIN@EUROHPC-JU.EUROPA.EU

BACKGROUND

- In January 2024 a package of measures was launched to support European startups and SMEs in the development of trustworthy artificial intelligence (AI) that respects EU values and rules (AI Factories).
- HPC(-AI) has a high entry barrier.
- Setup an action to lower the barrier and enable science, industry (SMEs, startups in particular) and public services to drive innovation.

HPC–AI SUPPORT CENTER (ASIC) CONCEPT

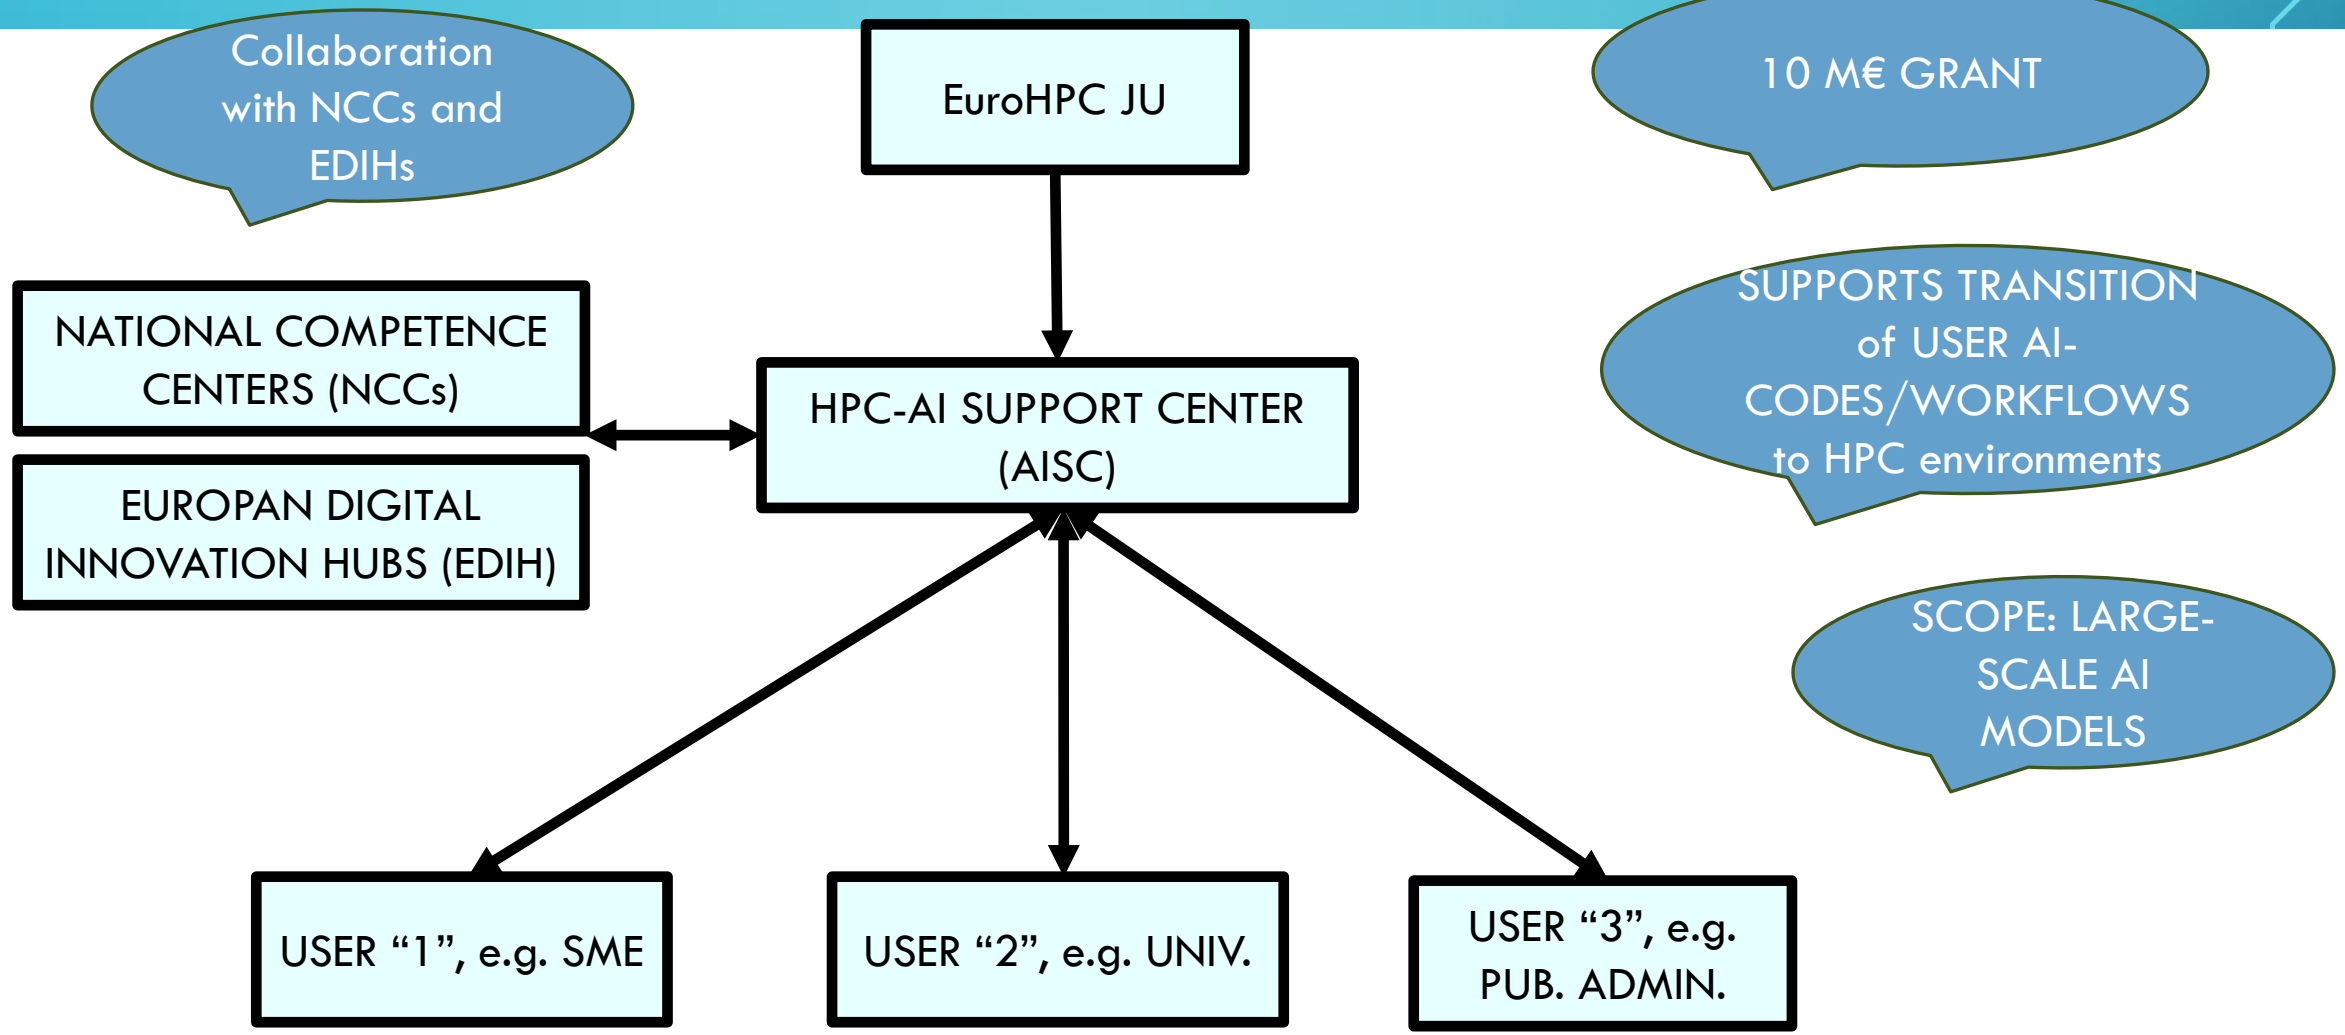

CURRENT SITUATION

- A consortium has been selected and invited for grant agreement preparation.
- Upcoming AI factories services.

THANKS!

For more information, feel free to visit our website and social media:

eurohpc-ju.europa.eu @euroHPC_JU eurohpc-ju @eurohpc-ju

EuroHPC User Day
User Forum Coordination Group

EuroHPC
Joint Undertaking

Introducing the EuroHPC User Forum

From the Multi-Annual Strategic Programme:

User Forum: EuroHPC User Forum will … promote **knowledge exchange**, **professional development**, and **collaboration** within the European HPC and Quantum communities. It shall be open, inclusive, independent, transparent, and responsive to the needs of its members. The Forum shall be made up of users from **academic**, **industrial**, and **public sector** sectors. It is an **open group** where discussions cover updates from EuroHPC, current developments in the European HPC community, collection of user requirements, and report on difficulties and issues they face using the EuroHPC JU infrastructure.

The creation of the EuroHPC User Forum shall have the following aim, support, and governance:

- The aim of the User Forum is to foster structured, coherent, and regular **communication and exchange** with all user communities and stakeholders;
- The mission and goals of the User Forum should be clear, concise, and relevant;
- The User Forum shall facilitate open consultation with user and scientific communities that also serves to highlight the EuroHPC vision;
- The User Forum shall have dedicated administrative support from the JU to ensure its sustainability and effectiveness;
- The JU shall establish a governance structure responsible for overseeing the Forum's activities and to collate and communicate **feedback on user requirements** to the Advisory Groups of EuroHPC as necessary;

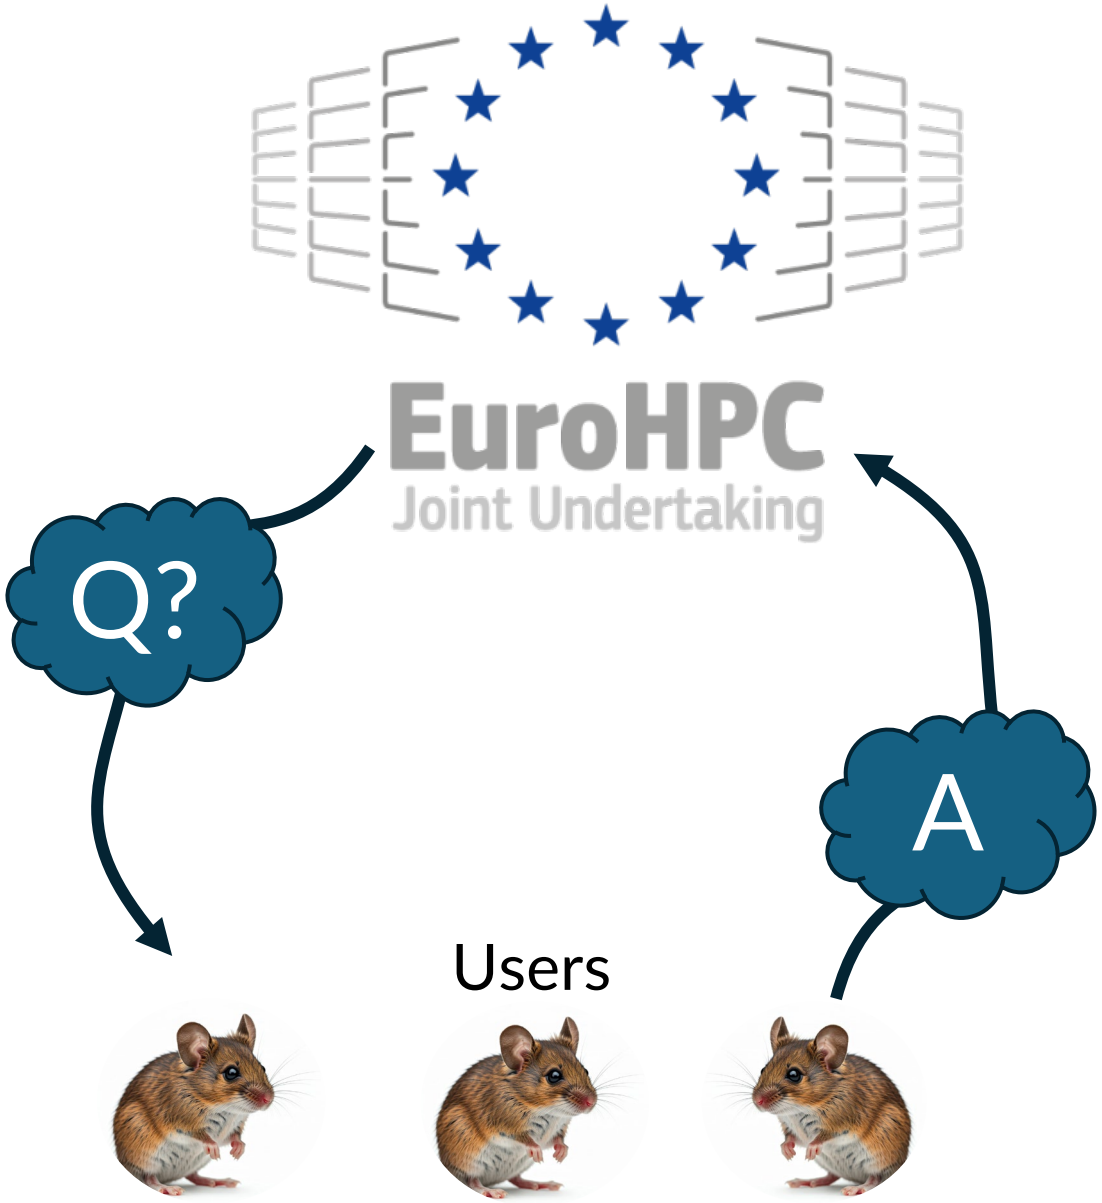
EuroHPC
Joint Undertaking
Q?
A
Users

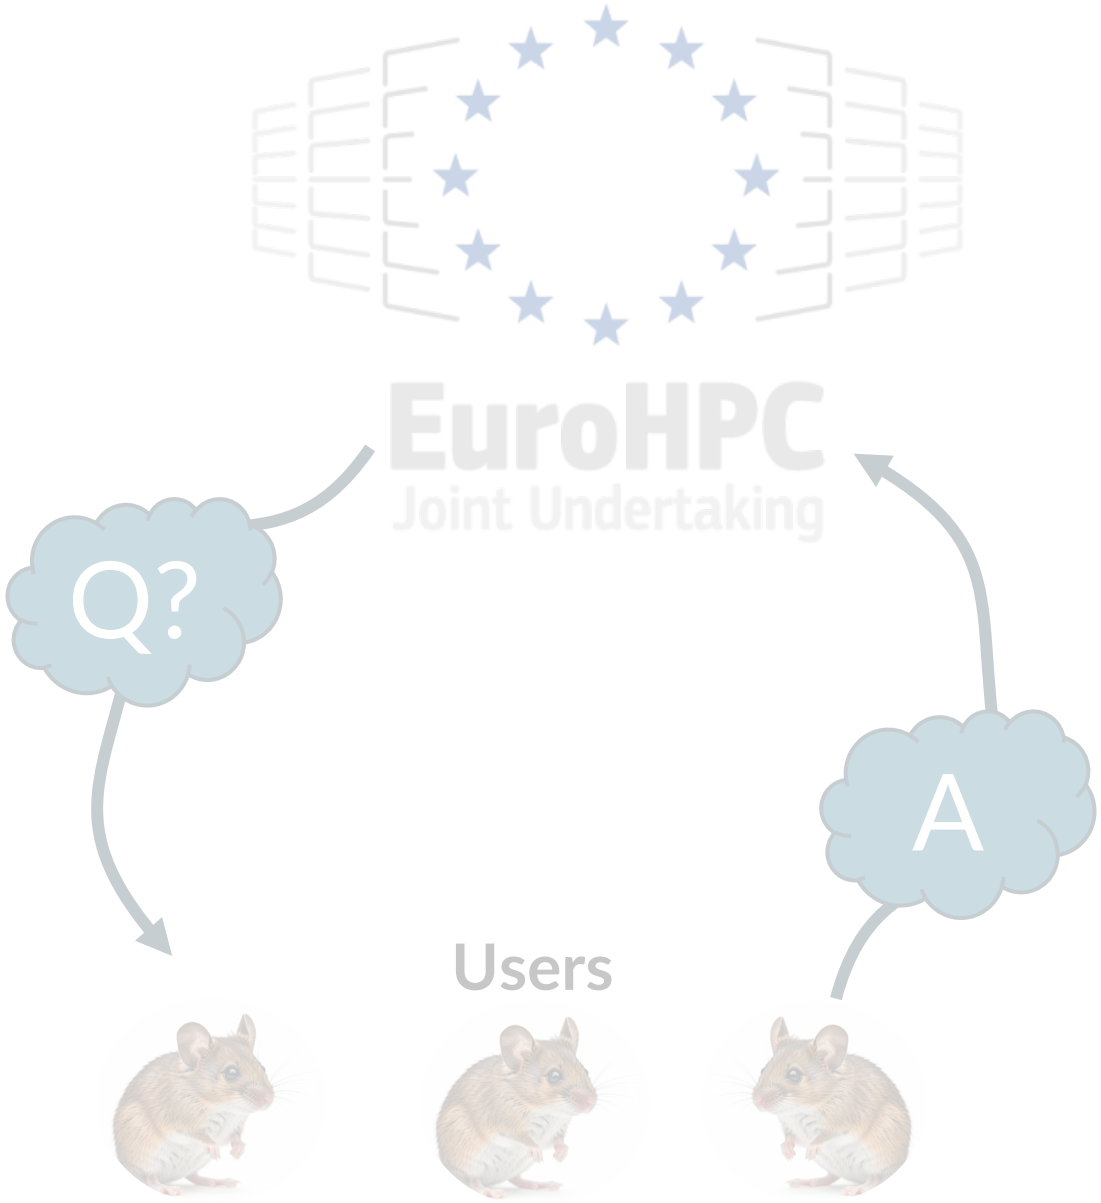
EuroHPC
Joint Undertaking
Q?
A
Users

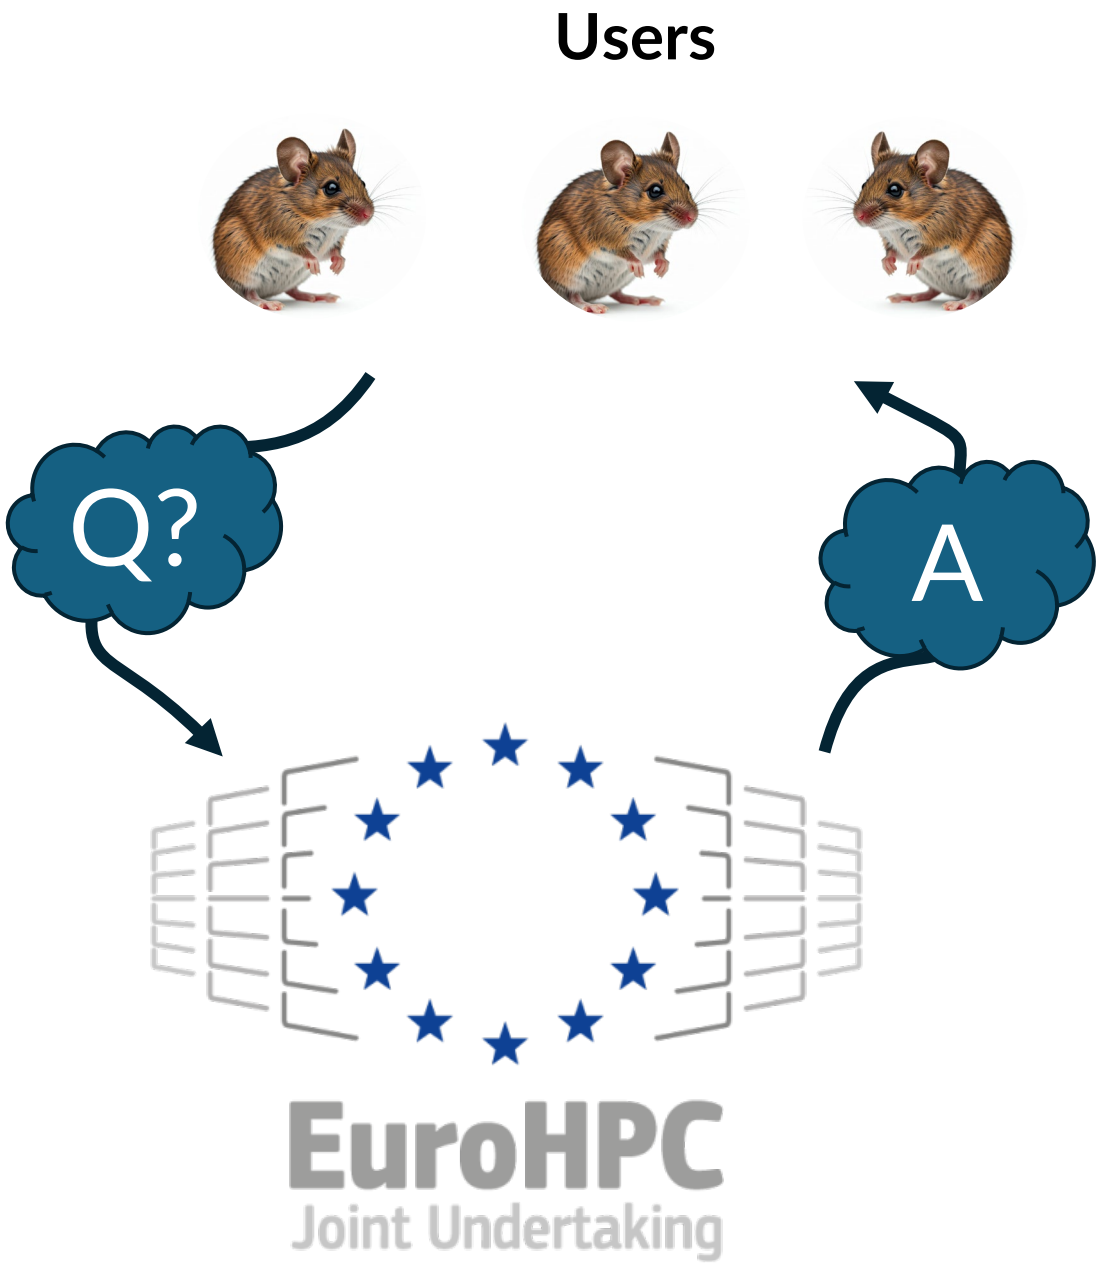
Users
Q?
A
EuroHPC
Joint Undertaking

User dialogue

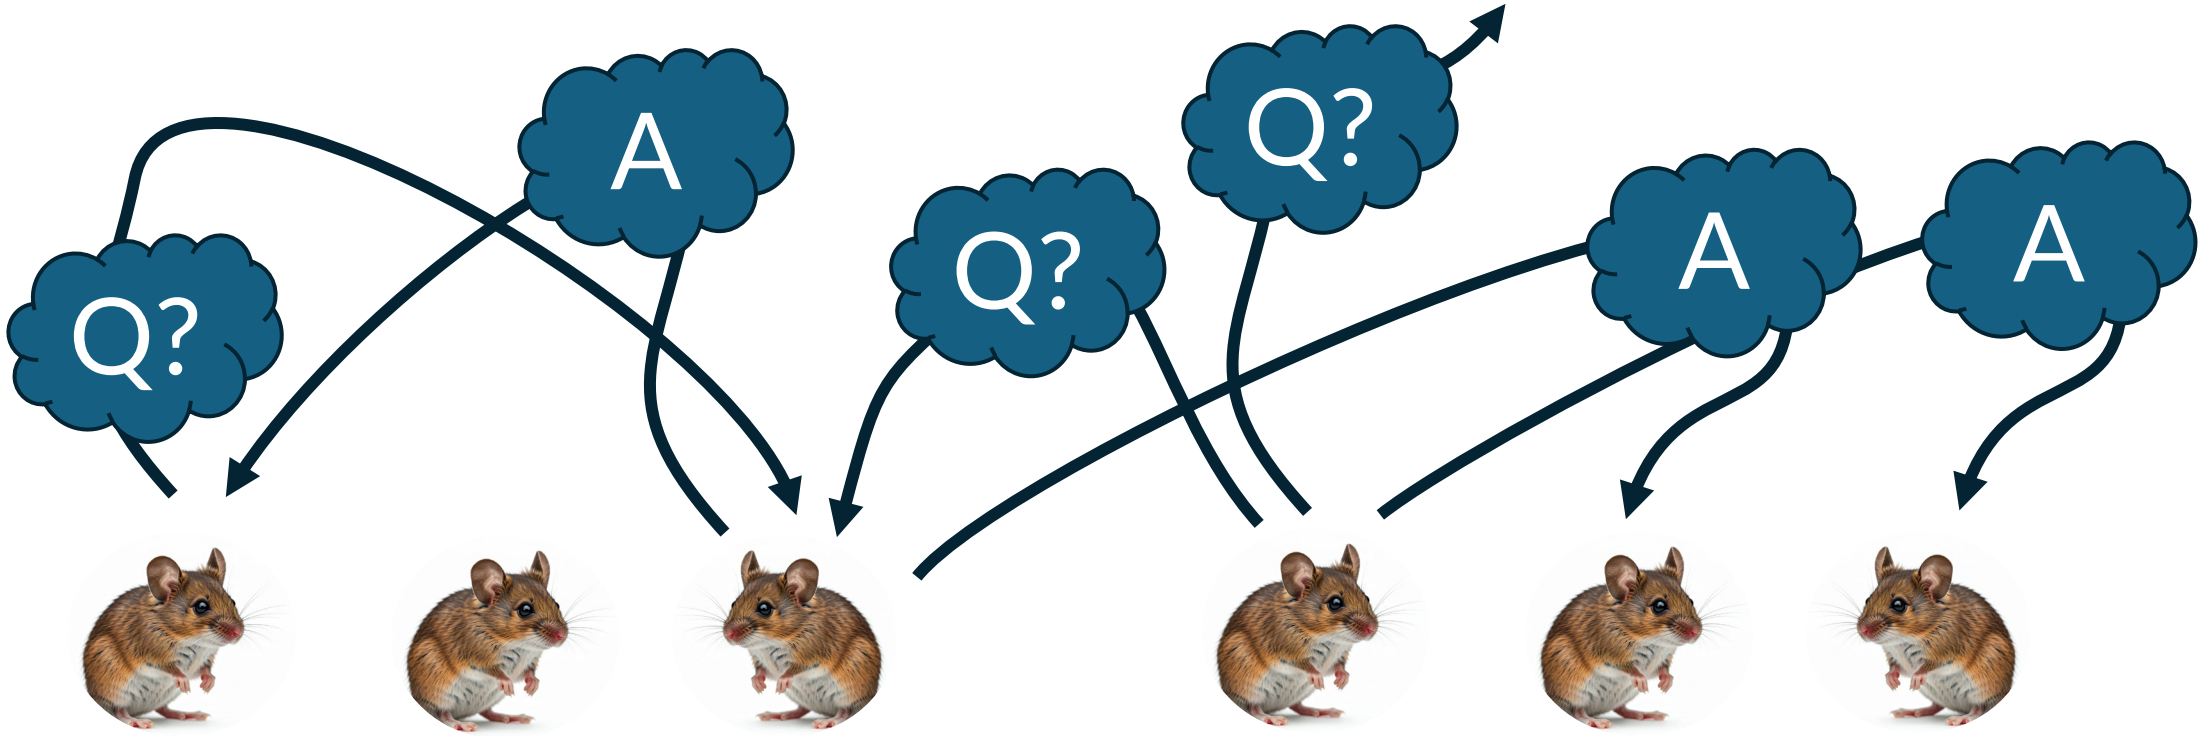

Who is involved in the Coordination Group

Chris Richardson (chair) Cambridge UK, Physics/Engineering finite element software development

Maria Girone (deputy chair) High Energy Physics, CERN-IT HPC Strategy Coordinator

Andrius Popovas (University of Oslo, Norway) HPC users for computational astrophysics

Xavier Besseron (University of Luxembourg) Research Scientist in Computer Science and Engineering

Ivan Carnimeo (CNR-IOM): Materials science/Chemistry/Physics developer of Quantum ESPRESSO, MAX (MAterials design at the eXascale) CoE, ICSC (Italian National Center for HPC)

Thomas Geenen (ECMWF, Destination Earth) - Climate and Weather Prediction and Earth System Digital Twin communities

Lara Peeters (VSC - HPC in Flanders, Belgium) - Digital Humanities

Sinéad Ryan – Lattice Quantum Chromodynamics TCD Dublin, Ireland

Zoe Cournia - Bioinformatics: Greece
Jean-Yves Verhaeghe - Computer science: France
Sergio Posada Perez - Nanoscienes: Belgium
Nuno Guerreiro - AI Neural network: Portugal
Peter Taborsky - AI LLM: Denmark
Matthias Meinke - Engineering: Germany

How to take part today!

- https://join.slack.com/t/eurohpcuserforum/shared_invite/zt-2r4ok2azn-mek1mzsAnH9z94okYbJr6A

slack

- Informal “chat” application
- Help to start conversations
- Not a “high security” platform

Discourse
may be next step

Reporting back to EuroHPC JU

Issues EuroHPC JU want to hear about…

How to encourage users to use resources efficiently

- penalties for users who do not?

AI usage needs

HPC Applications ecosystem needs:

- application software stack
- User environment (tools etc.)
- Compatibility across systems

January 2025

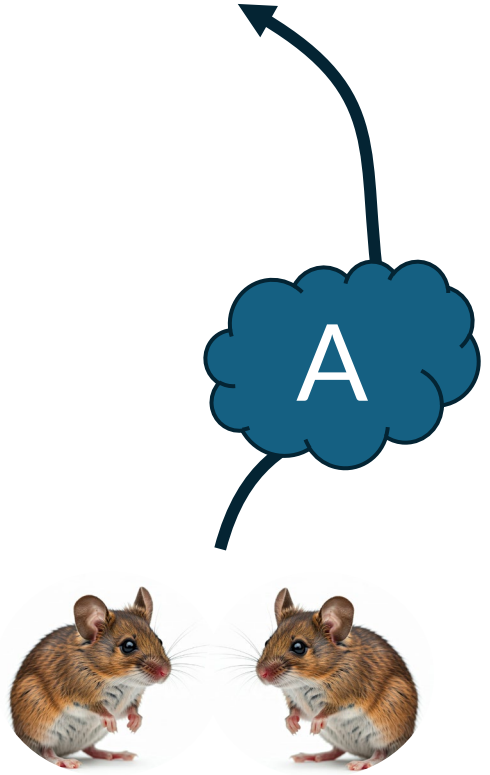

We need you!

- Please sign up to “slack” – or contact us by email
- The “user forum” will only work if you take part!

slack

Email: CEDeX-UFCG@eurohpc-ju.europa.eu

EuroHPC
Joint Undertaking