EUROHPC SUMMIT 2024

ANTWERP

TO EXASCALE
AND BEYOND

UNLEASHING THE POWER OF EUROPEAN HPC AND QUANTUM COMPUTING

EUROHPC SYSTEMS: CURRENT USAGE, ANTICIPATED NEEDS AND FUTURE CHALLENGES

EVANGELOS FLOROS
Head of Sector
Infrastructure
The EuroHPC Joint Undertaking

LILIT AXNER
Programme Manager
Infrastructure
The EuroHPC Joint Undertaking

DANIELA GALETTI
System, Storage and HPC Manager
CINECA

EMMANUEL ORY
Development Manager
CSC IT Center For
Science

BENEDIKT VON ST. VIETH
Division Head HPC, Cloud and Data Syst
Julich Supercomputing Centre

DAVID VICENTE
HPC User Support Manager
Barcelona Supercomputing
Center

STEPHANE REQUENA
CTO / Directeur Technique & Innovation
GENCI

EUROHPC SUMMIT 2024

ANTWERP

UNLEASHING THE POWER OF EUROPEAN HPC AND QUANTUM COMPUTING

European High Performance Computing Joint Undertaking

Usage Statistics of the EuroHPC JU Systems for the Year 2023

Dr. Lilit Axner
Programme Officer Infrastructure

EUROHPC SYSTEMS
2019 → 2023

PRE-EXASCALE

PETASCALE

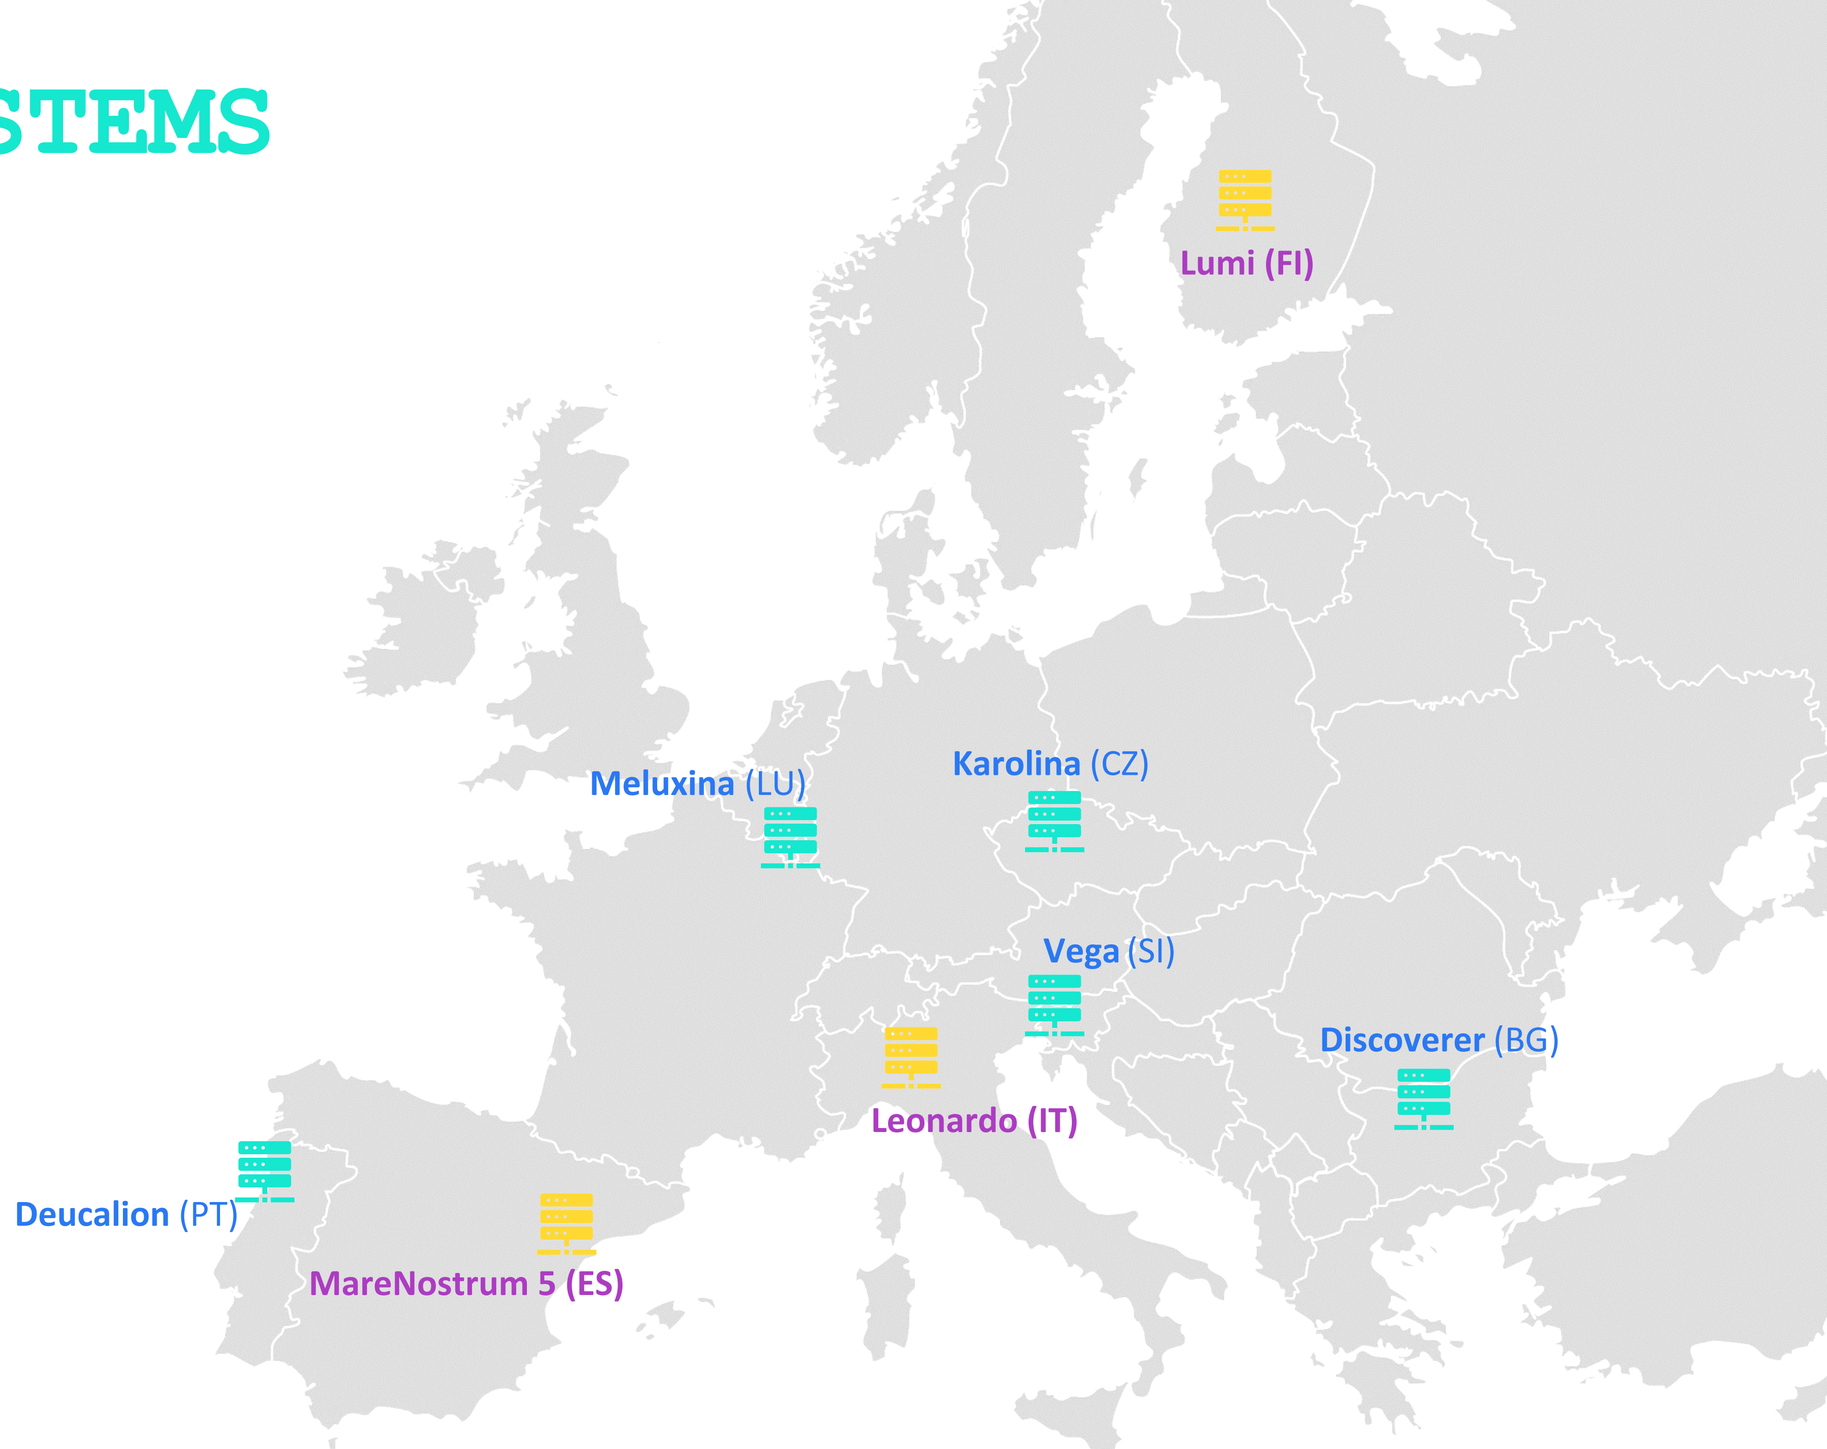

EUROHPC SYSTEMS
2024 → 2026

EXASCALE

PRE-EXASCALE

PETASCALE

MID-RANGE

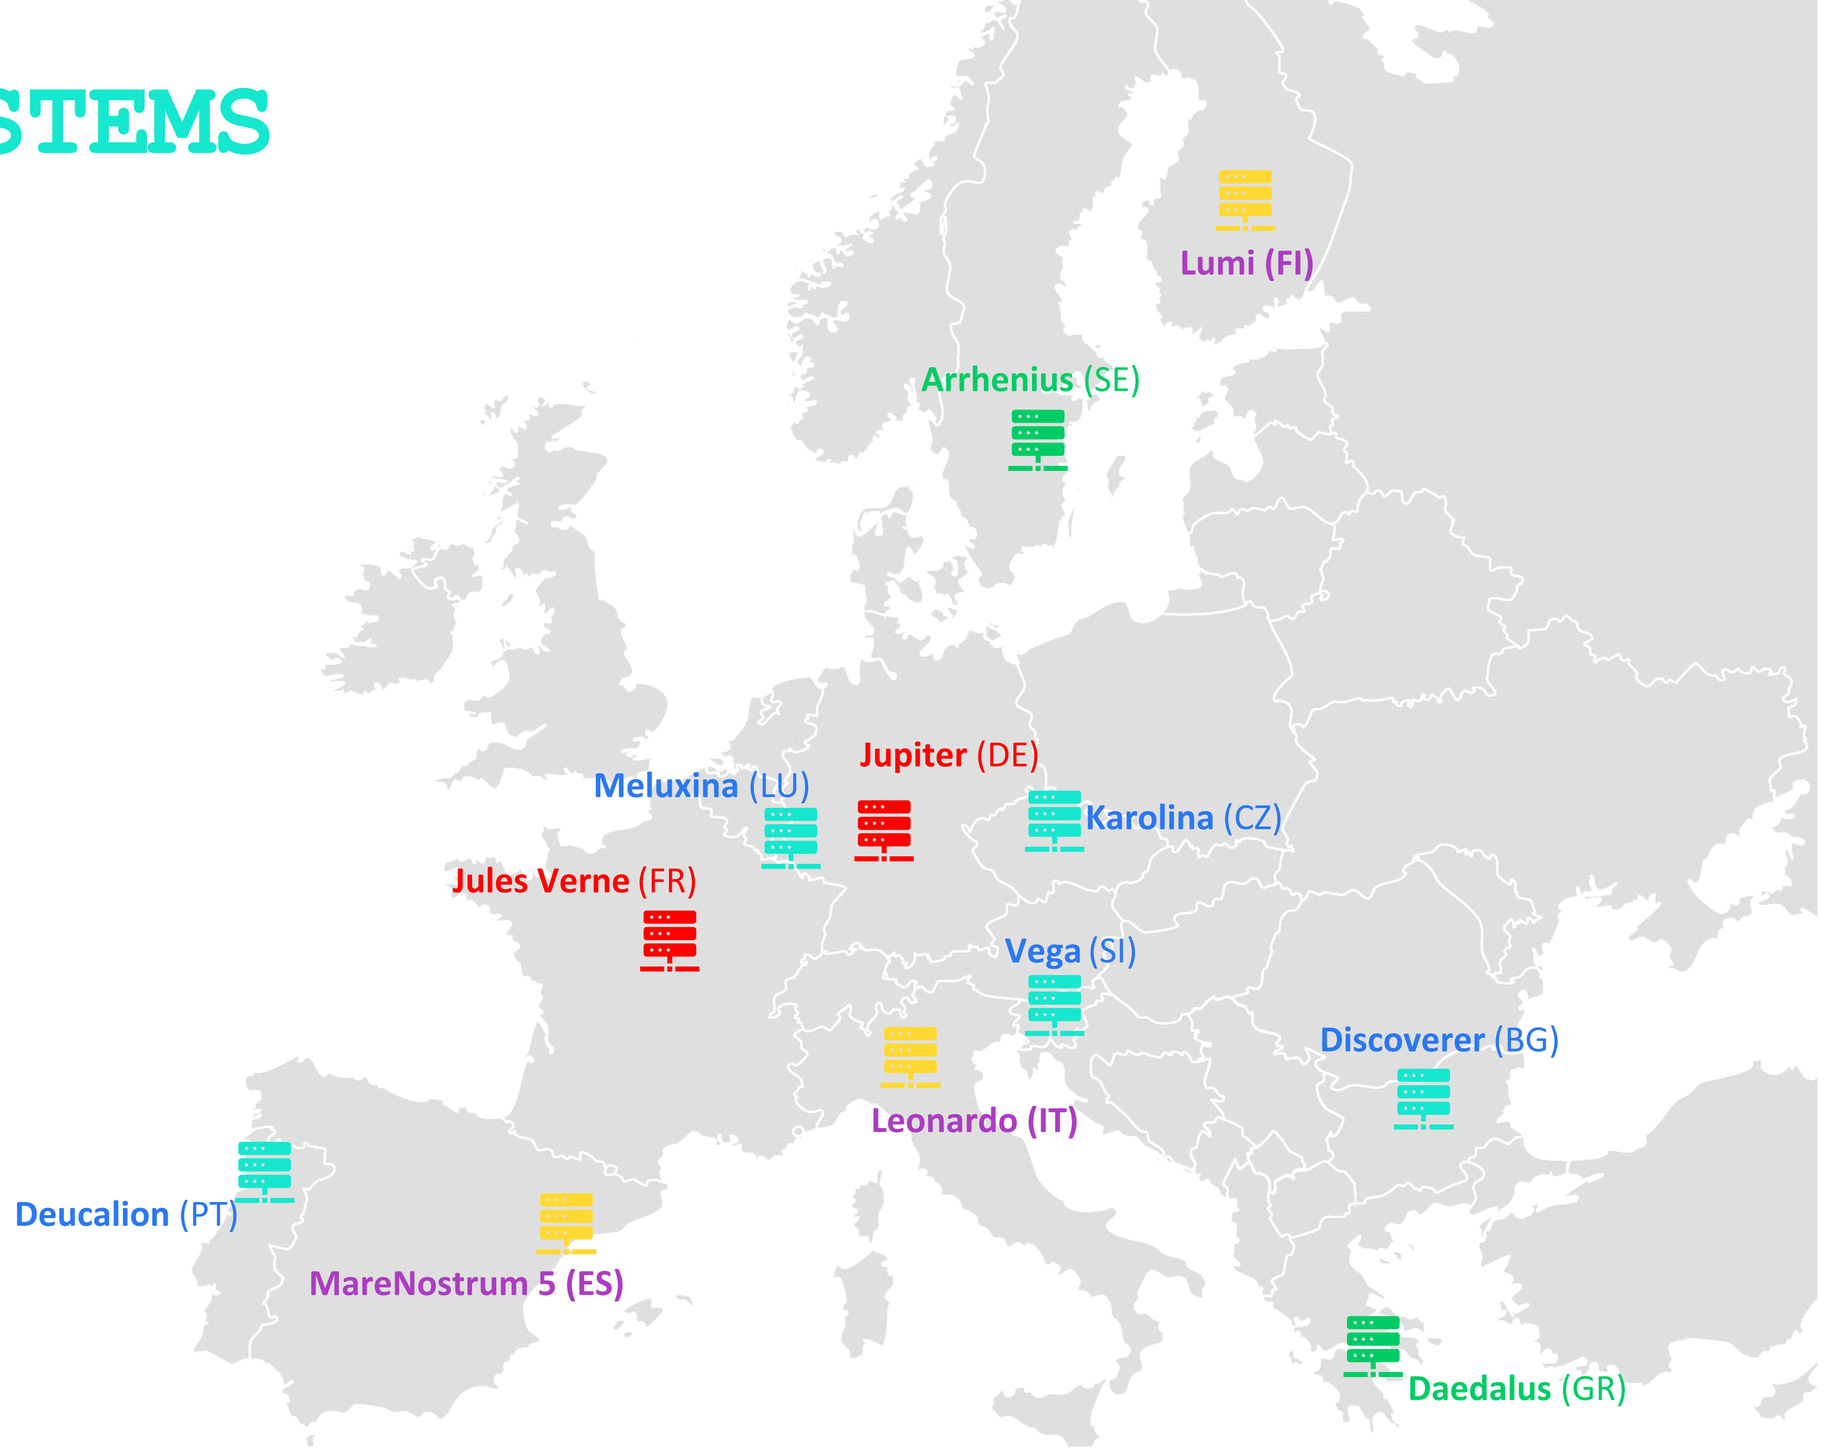

596 projects started their executions on EuroHPC JU systems in 2023 (Distribution per Country of the PIs)

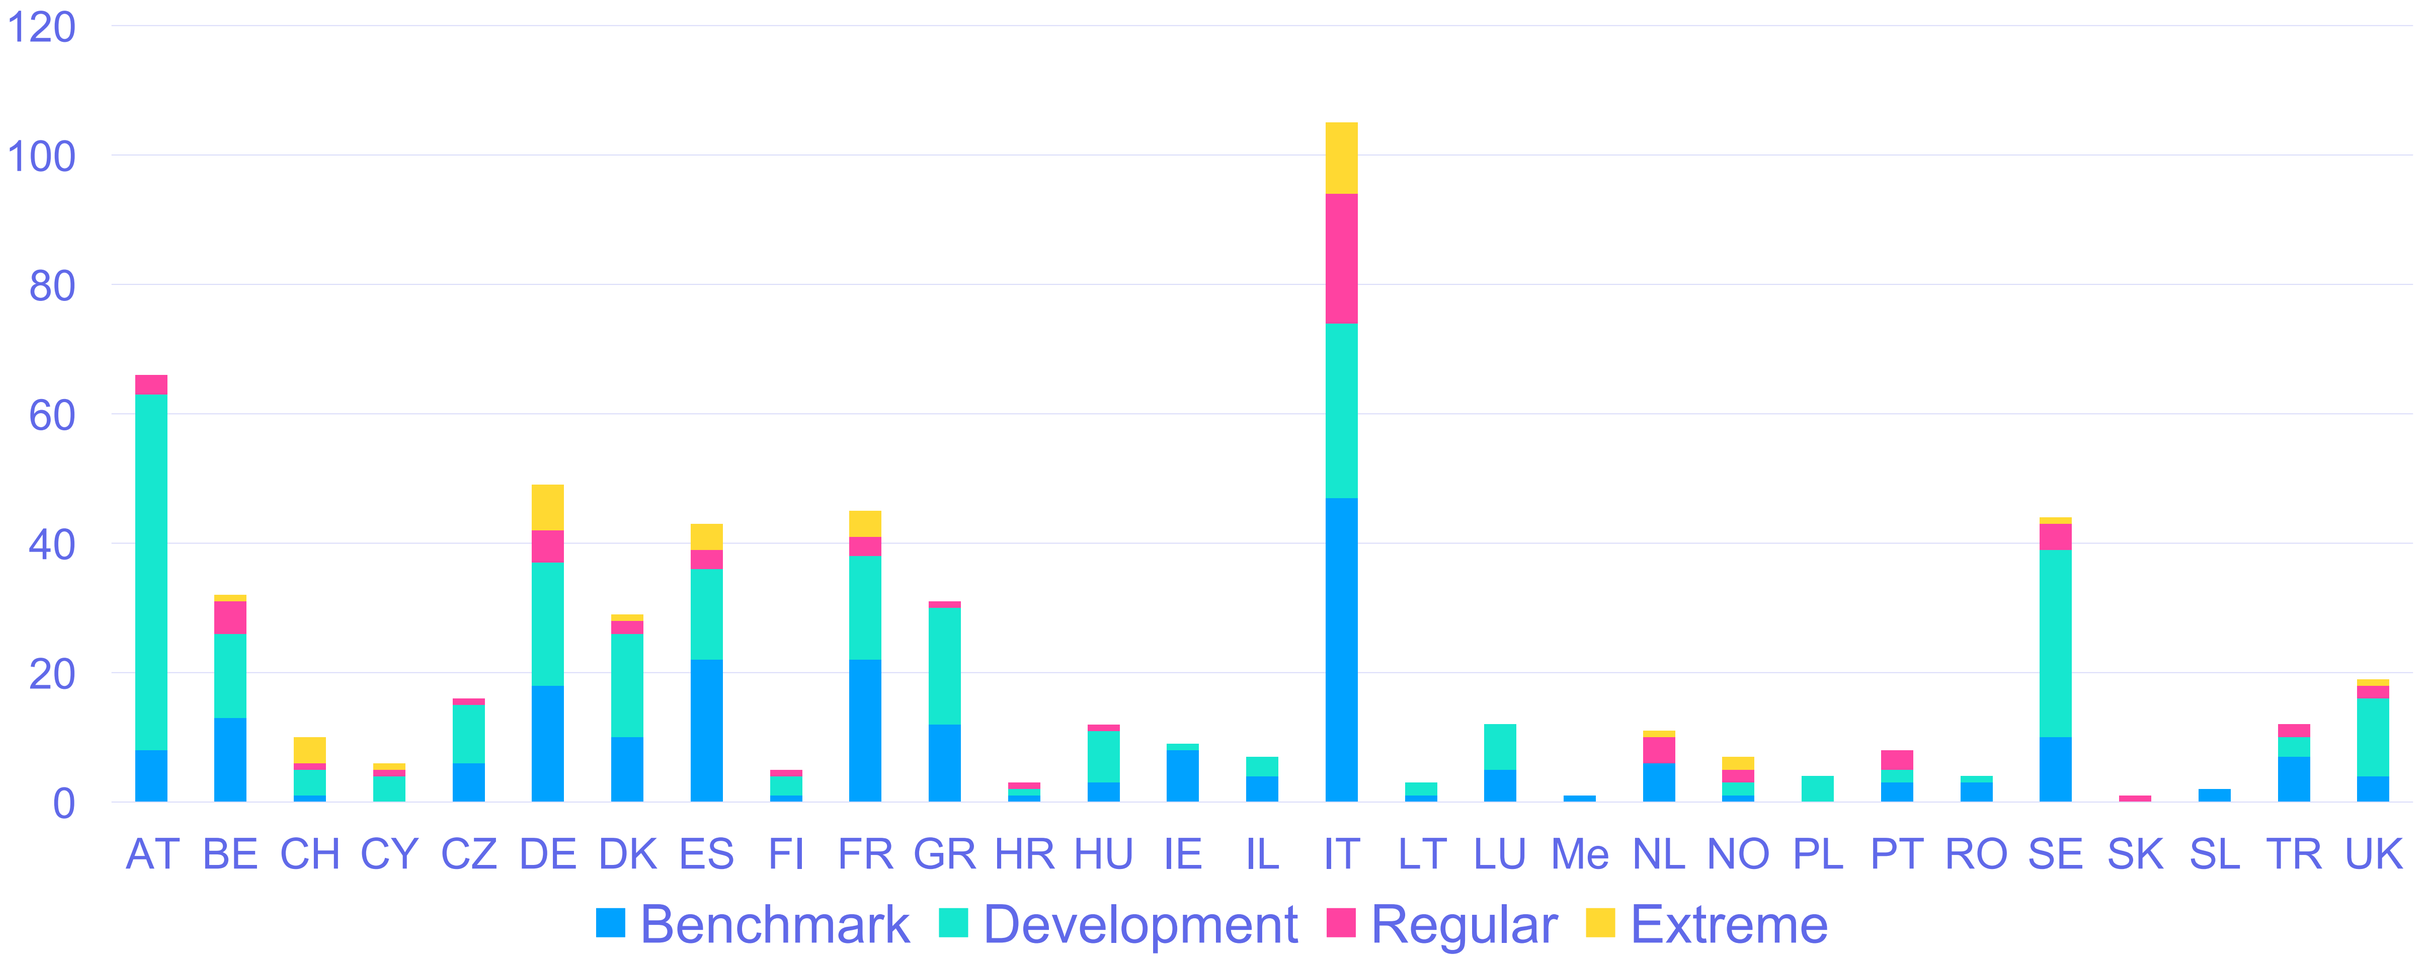

167 AI projects started their executions on EuroHPC JU systems in 2023 (Distribution per Country of the PIs, per type of applicant)

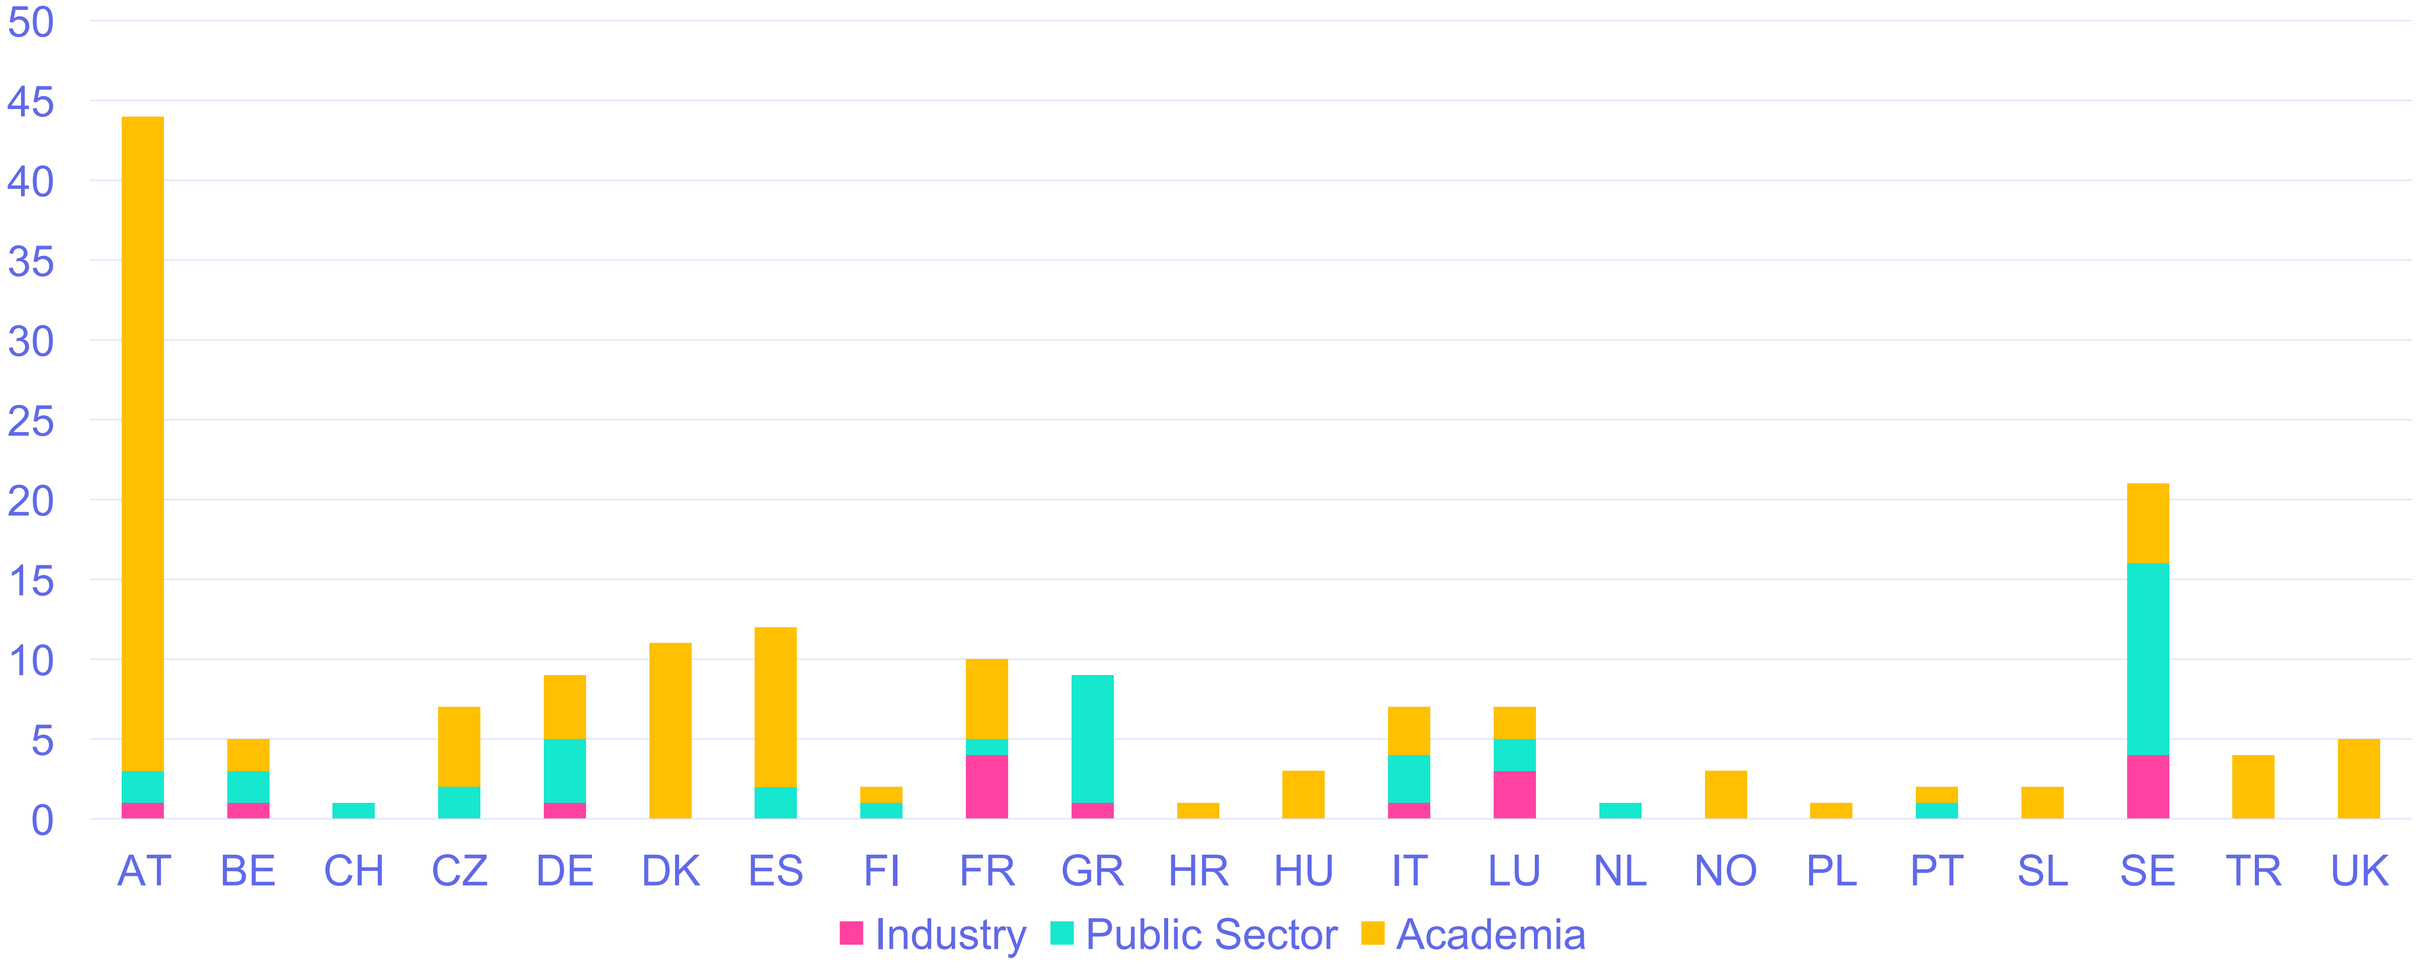

167 active AI projects distribution per system per type of applicant

94% of projects were applied through Benchmark and Development calls

10% of AI projects are from Industry

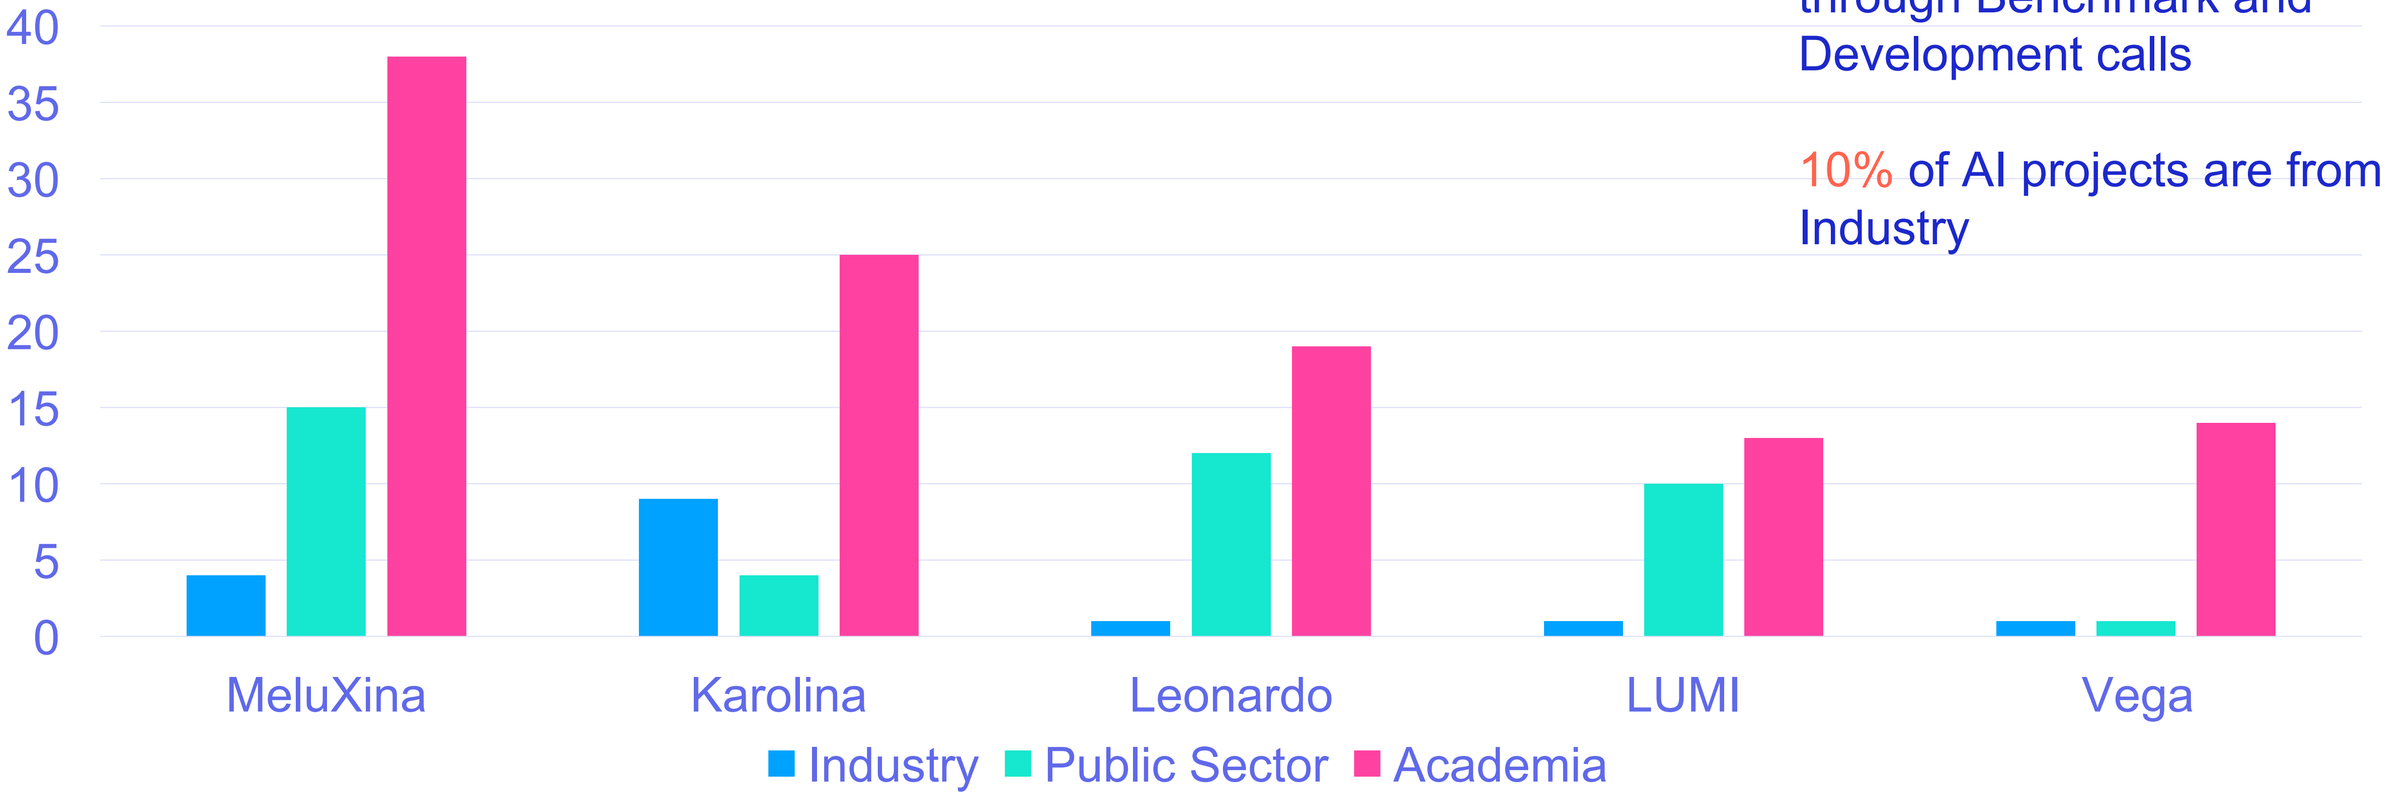

(*) Leonardo was available for access since June 2023
(**) LUMI-G was on maintenances few times during the 2023 as it was newly installed

596 active projects distribution per system per type of EuroHPC JU access modes

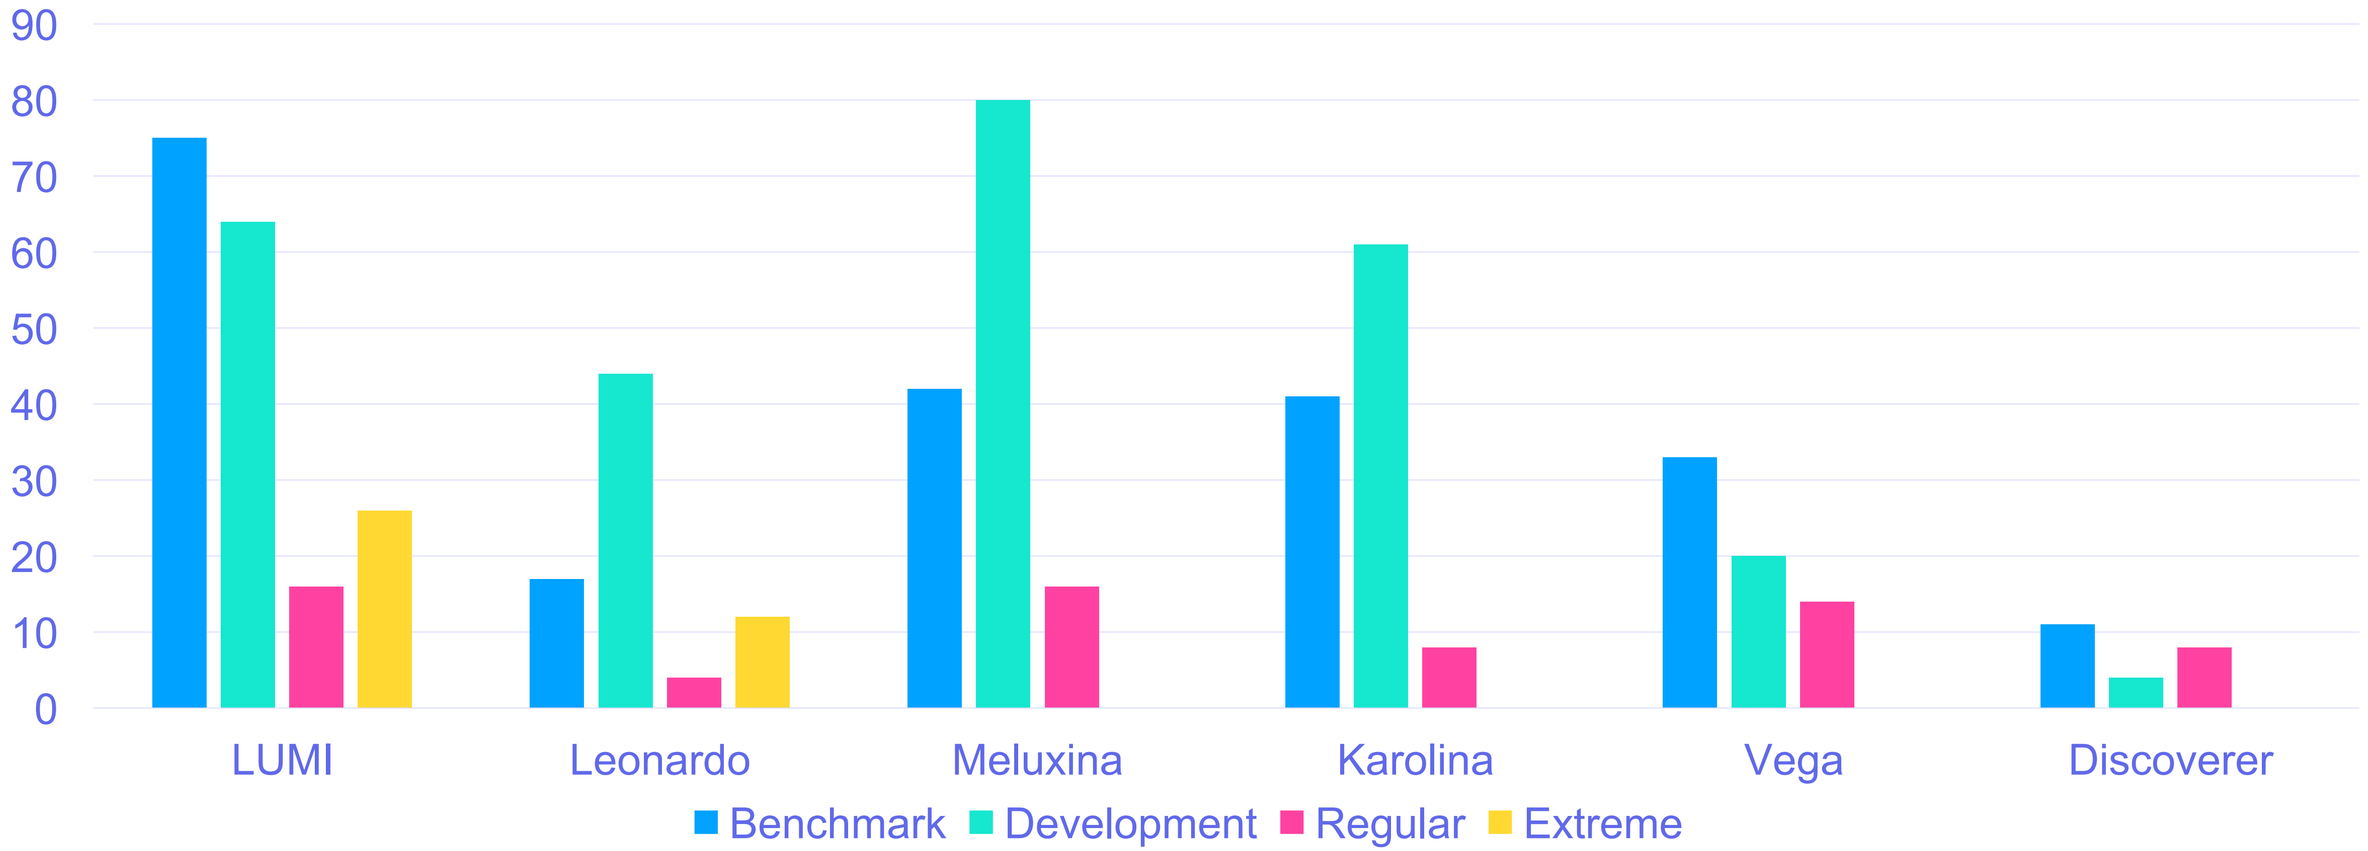

Gender ratio of 596 projects

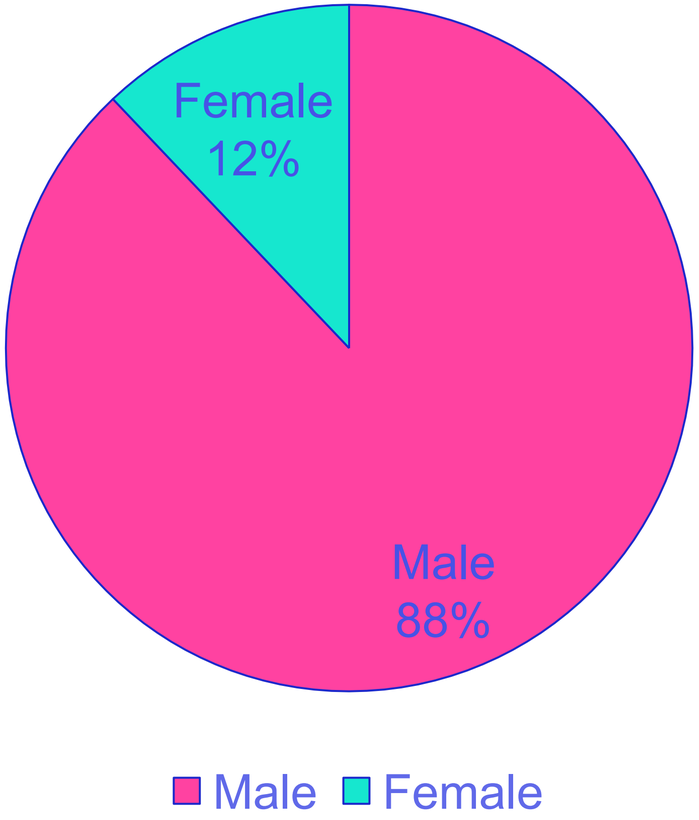

596 projects distribution per type of applicant

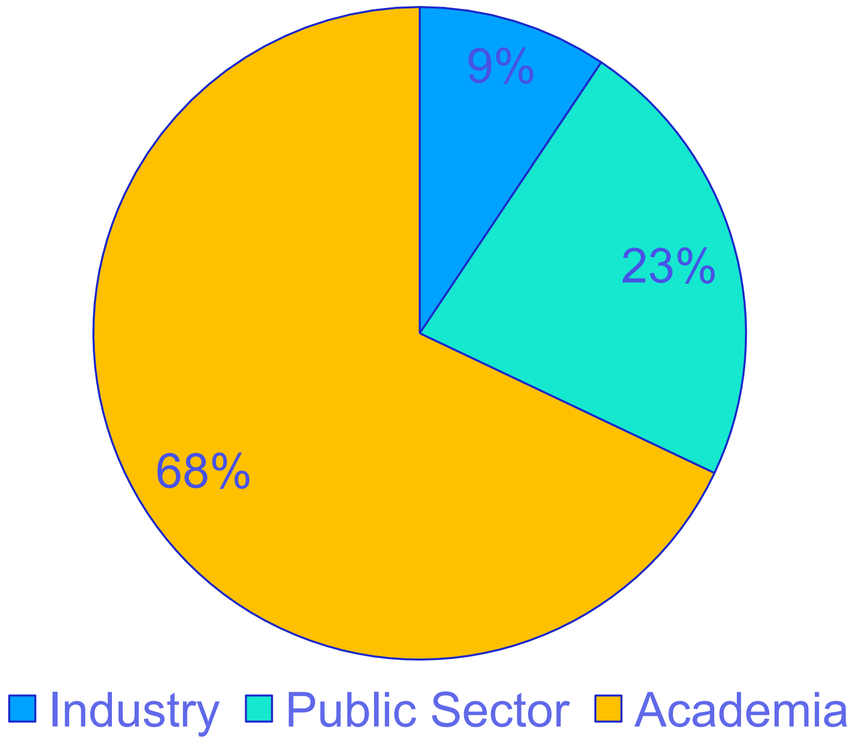

Resources that were used by 596 projects during 2023 per EuroHPC JU system

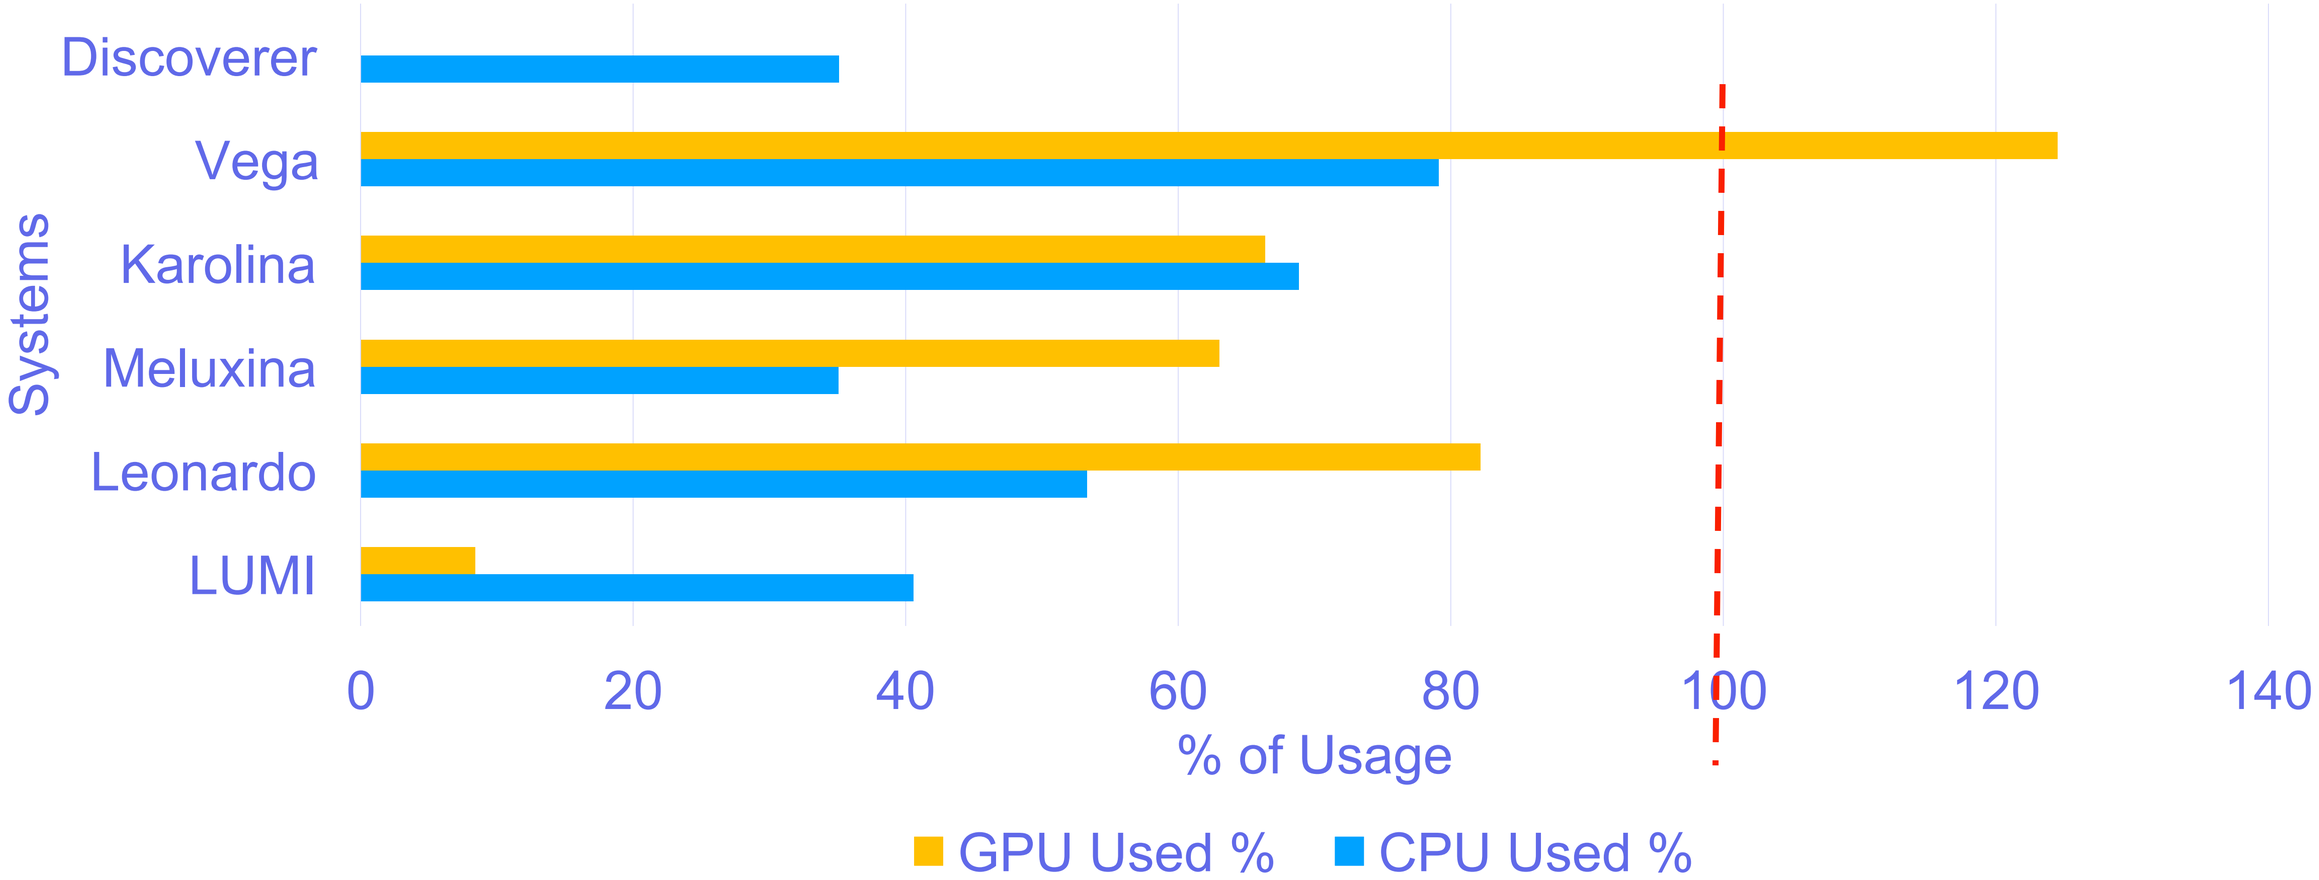

(*) Leonardo was available for access since June 2023

(**) LUMI-G was on maintenances few times during the 2023 as it was newly installed. (***) Discoverer has only CPU partition.

Access for AI and Data-Intensive Applications

AI Access: Start to use after 1 month, for the duration of **1** year

Resources: Up to **35 000GPU** Node hours

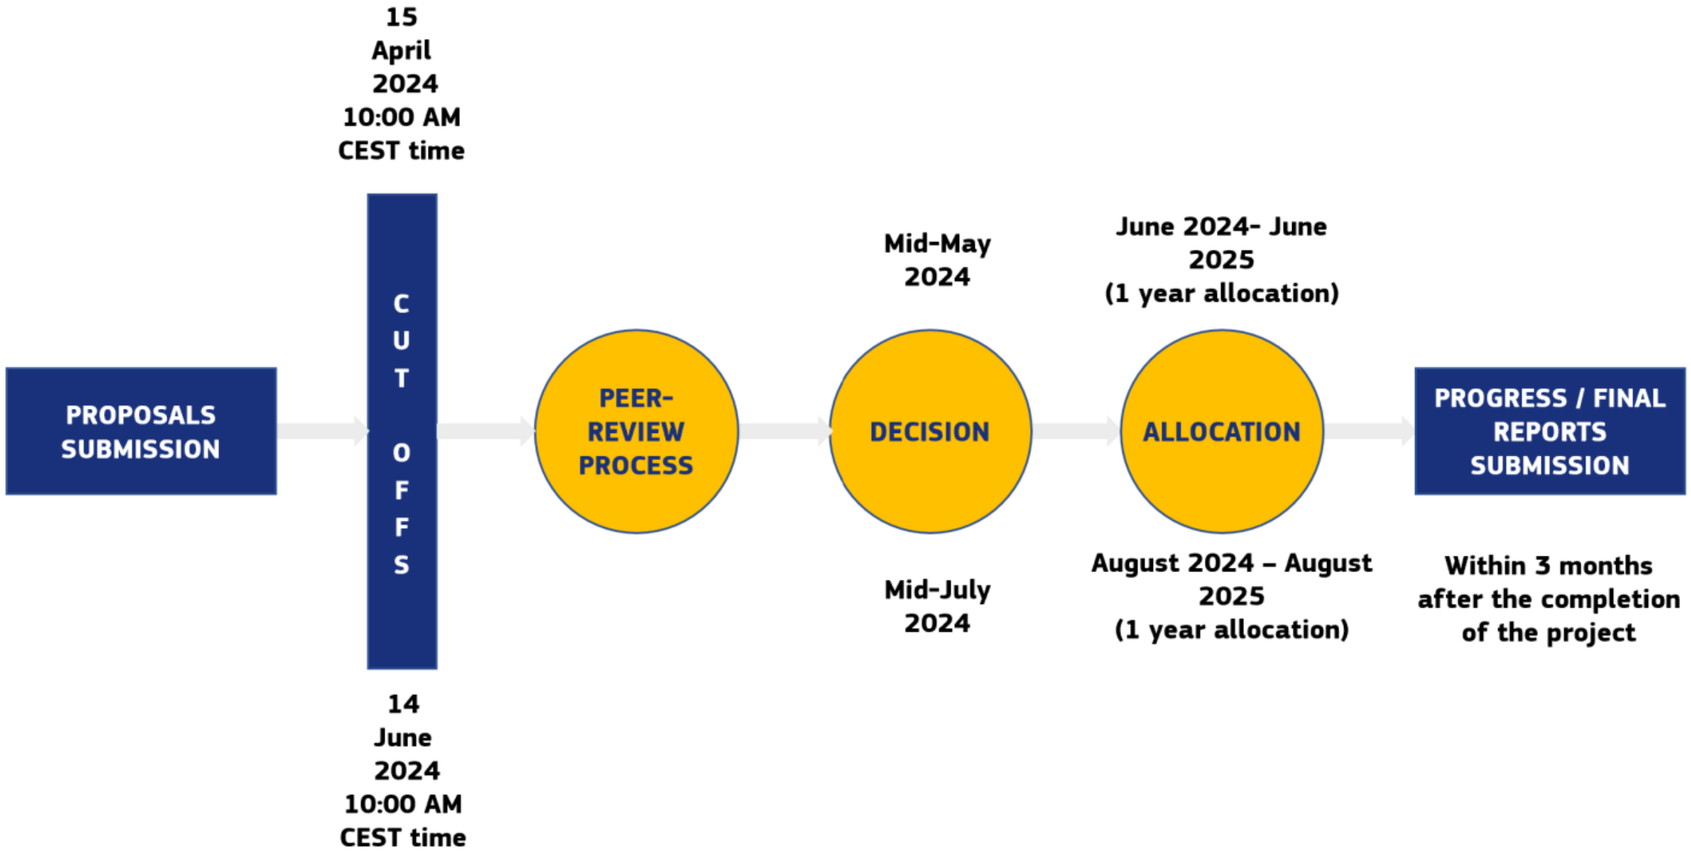

EPICURE: Software support for the EuroHPC JU successful applicants

- 14 HPC centres across Europe
- 20 HPC software experts
- 988 month of experts' work on **European** HPC and AI software of the **EuroHPC JU successful applicants** during 4 years
- **Port, Scale, Optimise, Benchmark** = Keywords of the project
- Organise at least 24 HPC and AI training events and hackathons
- Create a one-stop-shop portal for users
- Produce at least 15 Best Practice Guides and similar manuals for users

Future Needs….

- Already in progress: Substantial user software support (Epicure, CoEs, etc…) for the move from pre-exascale to Exascale
- Already in progress: Federated user-friendly EuroHPC ecosystem

- Increase in demand of Computing Resources
- Demand for enhanced and robust interconnectivity of resources (**Hyperdata** hypermobility = Hyperconnectivity)
- Increased **diversity** of “non-traditional” HPC users
- Increase in demand of **diversity** of HPC software supporting activities
- Increase in demand of **diversity** of HPC middleware supporting activities

EUROHPC SUMMIT 2024

ANTWERP

TO EXASCALE
AND BEYOND

UNLEASHING THE POWER OF EUROPEAN HPC AND QUANTUM COMPUTING

THANK YOU!

EuroHPC
Joint Undertaking

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

eurohpc-ju.europa.eu

@euroHPC_JU

eurohpc-ju

EUROHPC SUMMIT 2024

ANTWERP

UNLEASHING THE POWER OF EUROPEAN HPC AND QUANTUM COMPUTING

EuroHPC Systems: current usage, anticipated needs and future challenges

LUMI (Large Unified Modern Infrastructure)

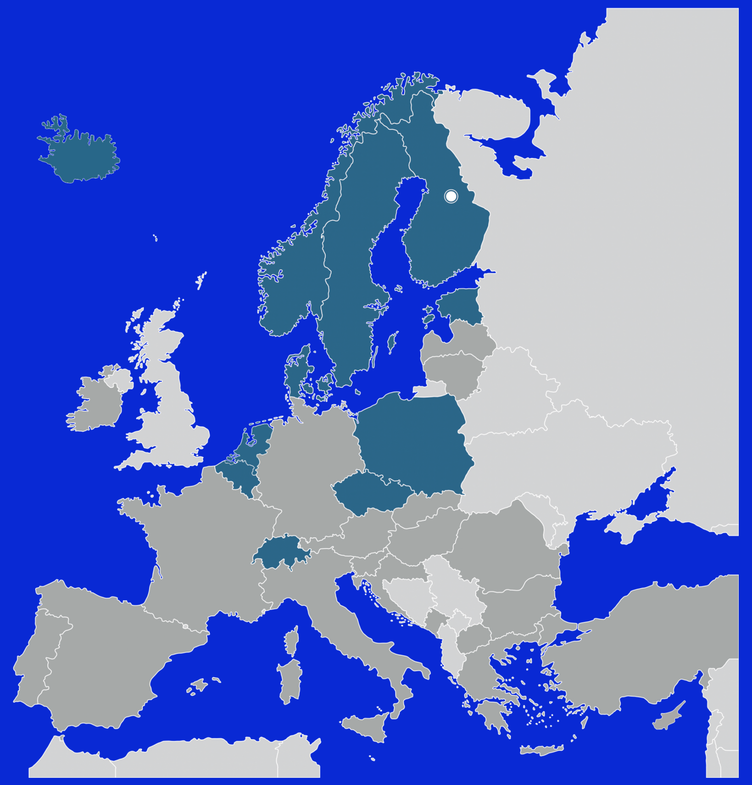

emmanuel.ory@csc.fi - EuroHPC Summit – Antwerp 18-21 March 2024

Current usage

LUMI-C in production since December 2021

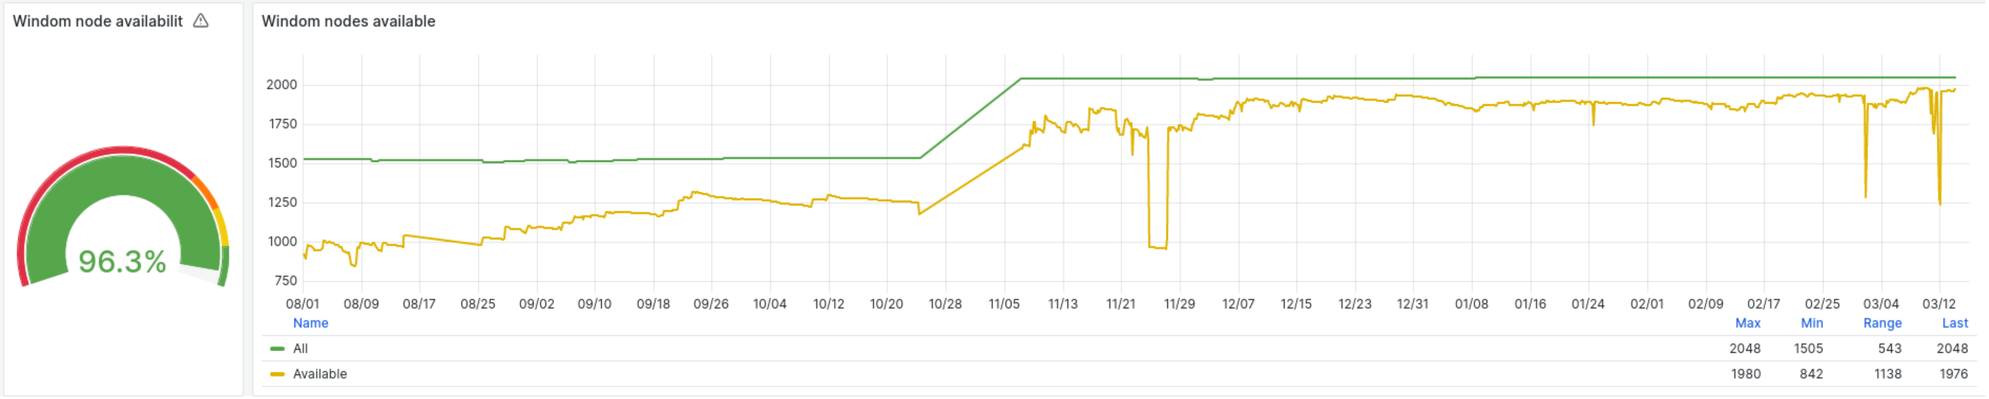

LUMI-G in production since November 2022

Current usage

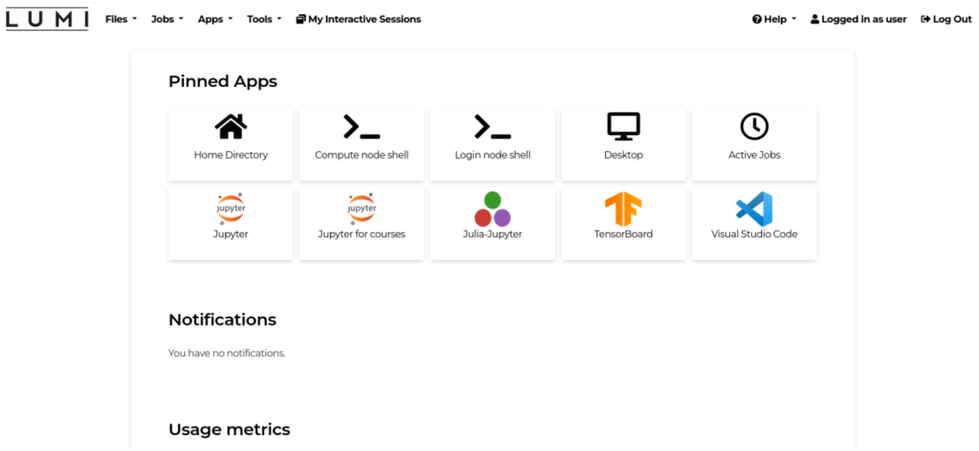

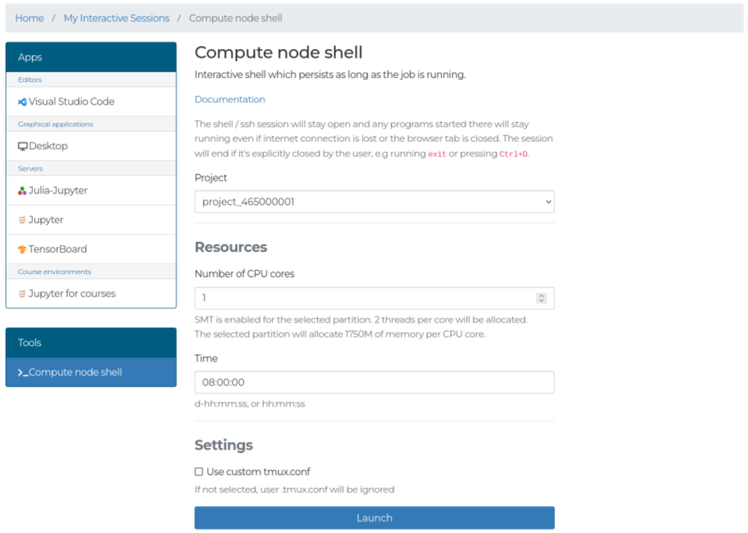

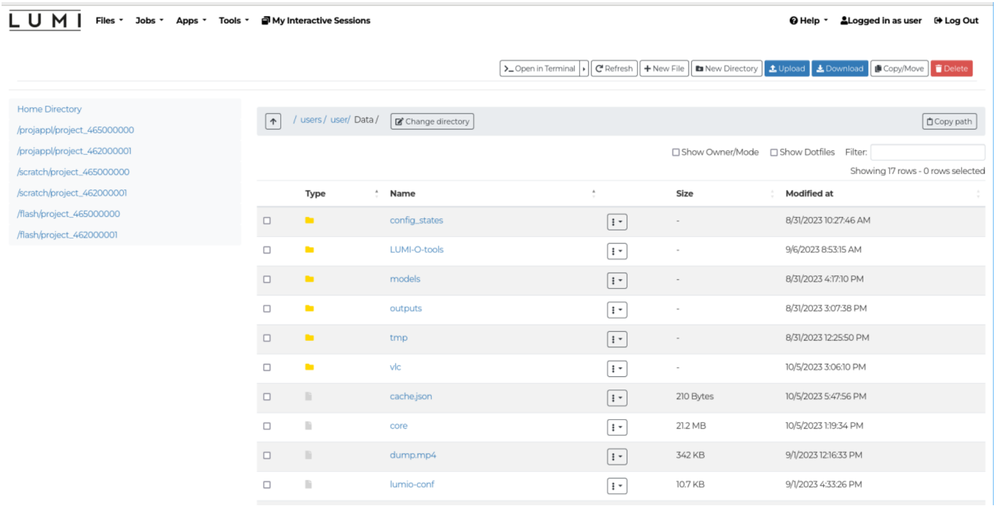

- ~1500 projects (incl. ~500 Finnish projects & 50 commercial projects)
- LUMI interface: ~70 unique users per day & ~3500 jobs launched in 2024

Current usage

- LUMI is also connected to **Helmi quantum computer** (VTT Technical Research Centre of Finland)
  - 263525 Quantum Processing Unit seconds in 2023!

- **Full machine runs** https://docs.lumi-supercomputer.eu/runjobs/scheduled-jobs/hero-runs/
- **Recent use cases:**
  - Accelerating Materials Discovery and Design with Machine Learning (EuroCC Belgium)
  - LUMI powers the study of light scattering in space (CSC)

Anticipated needs and future challenges

- LUMI RT ticketing system: ~4000 tickets
  - 1st pre-exascale system in Europe
  - Adapting the skills & competences to better support the users
  - Managing users/projects from 11 consortium countries + EuroHPC JU

- **250M€ secured by the Finnish government for LUMI-next**

EUROHPC SUMMIT 2024

ANTWERP

MareNostrum 5

EuroHPC Systems: Current Usage, Anticipated Needs and Future Challenges

David Vicente – BSC - CNS

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

UNLEASHING THE POWER OF EUROPEAN HPC AND QUANTUM COMPUTING

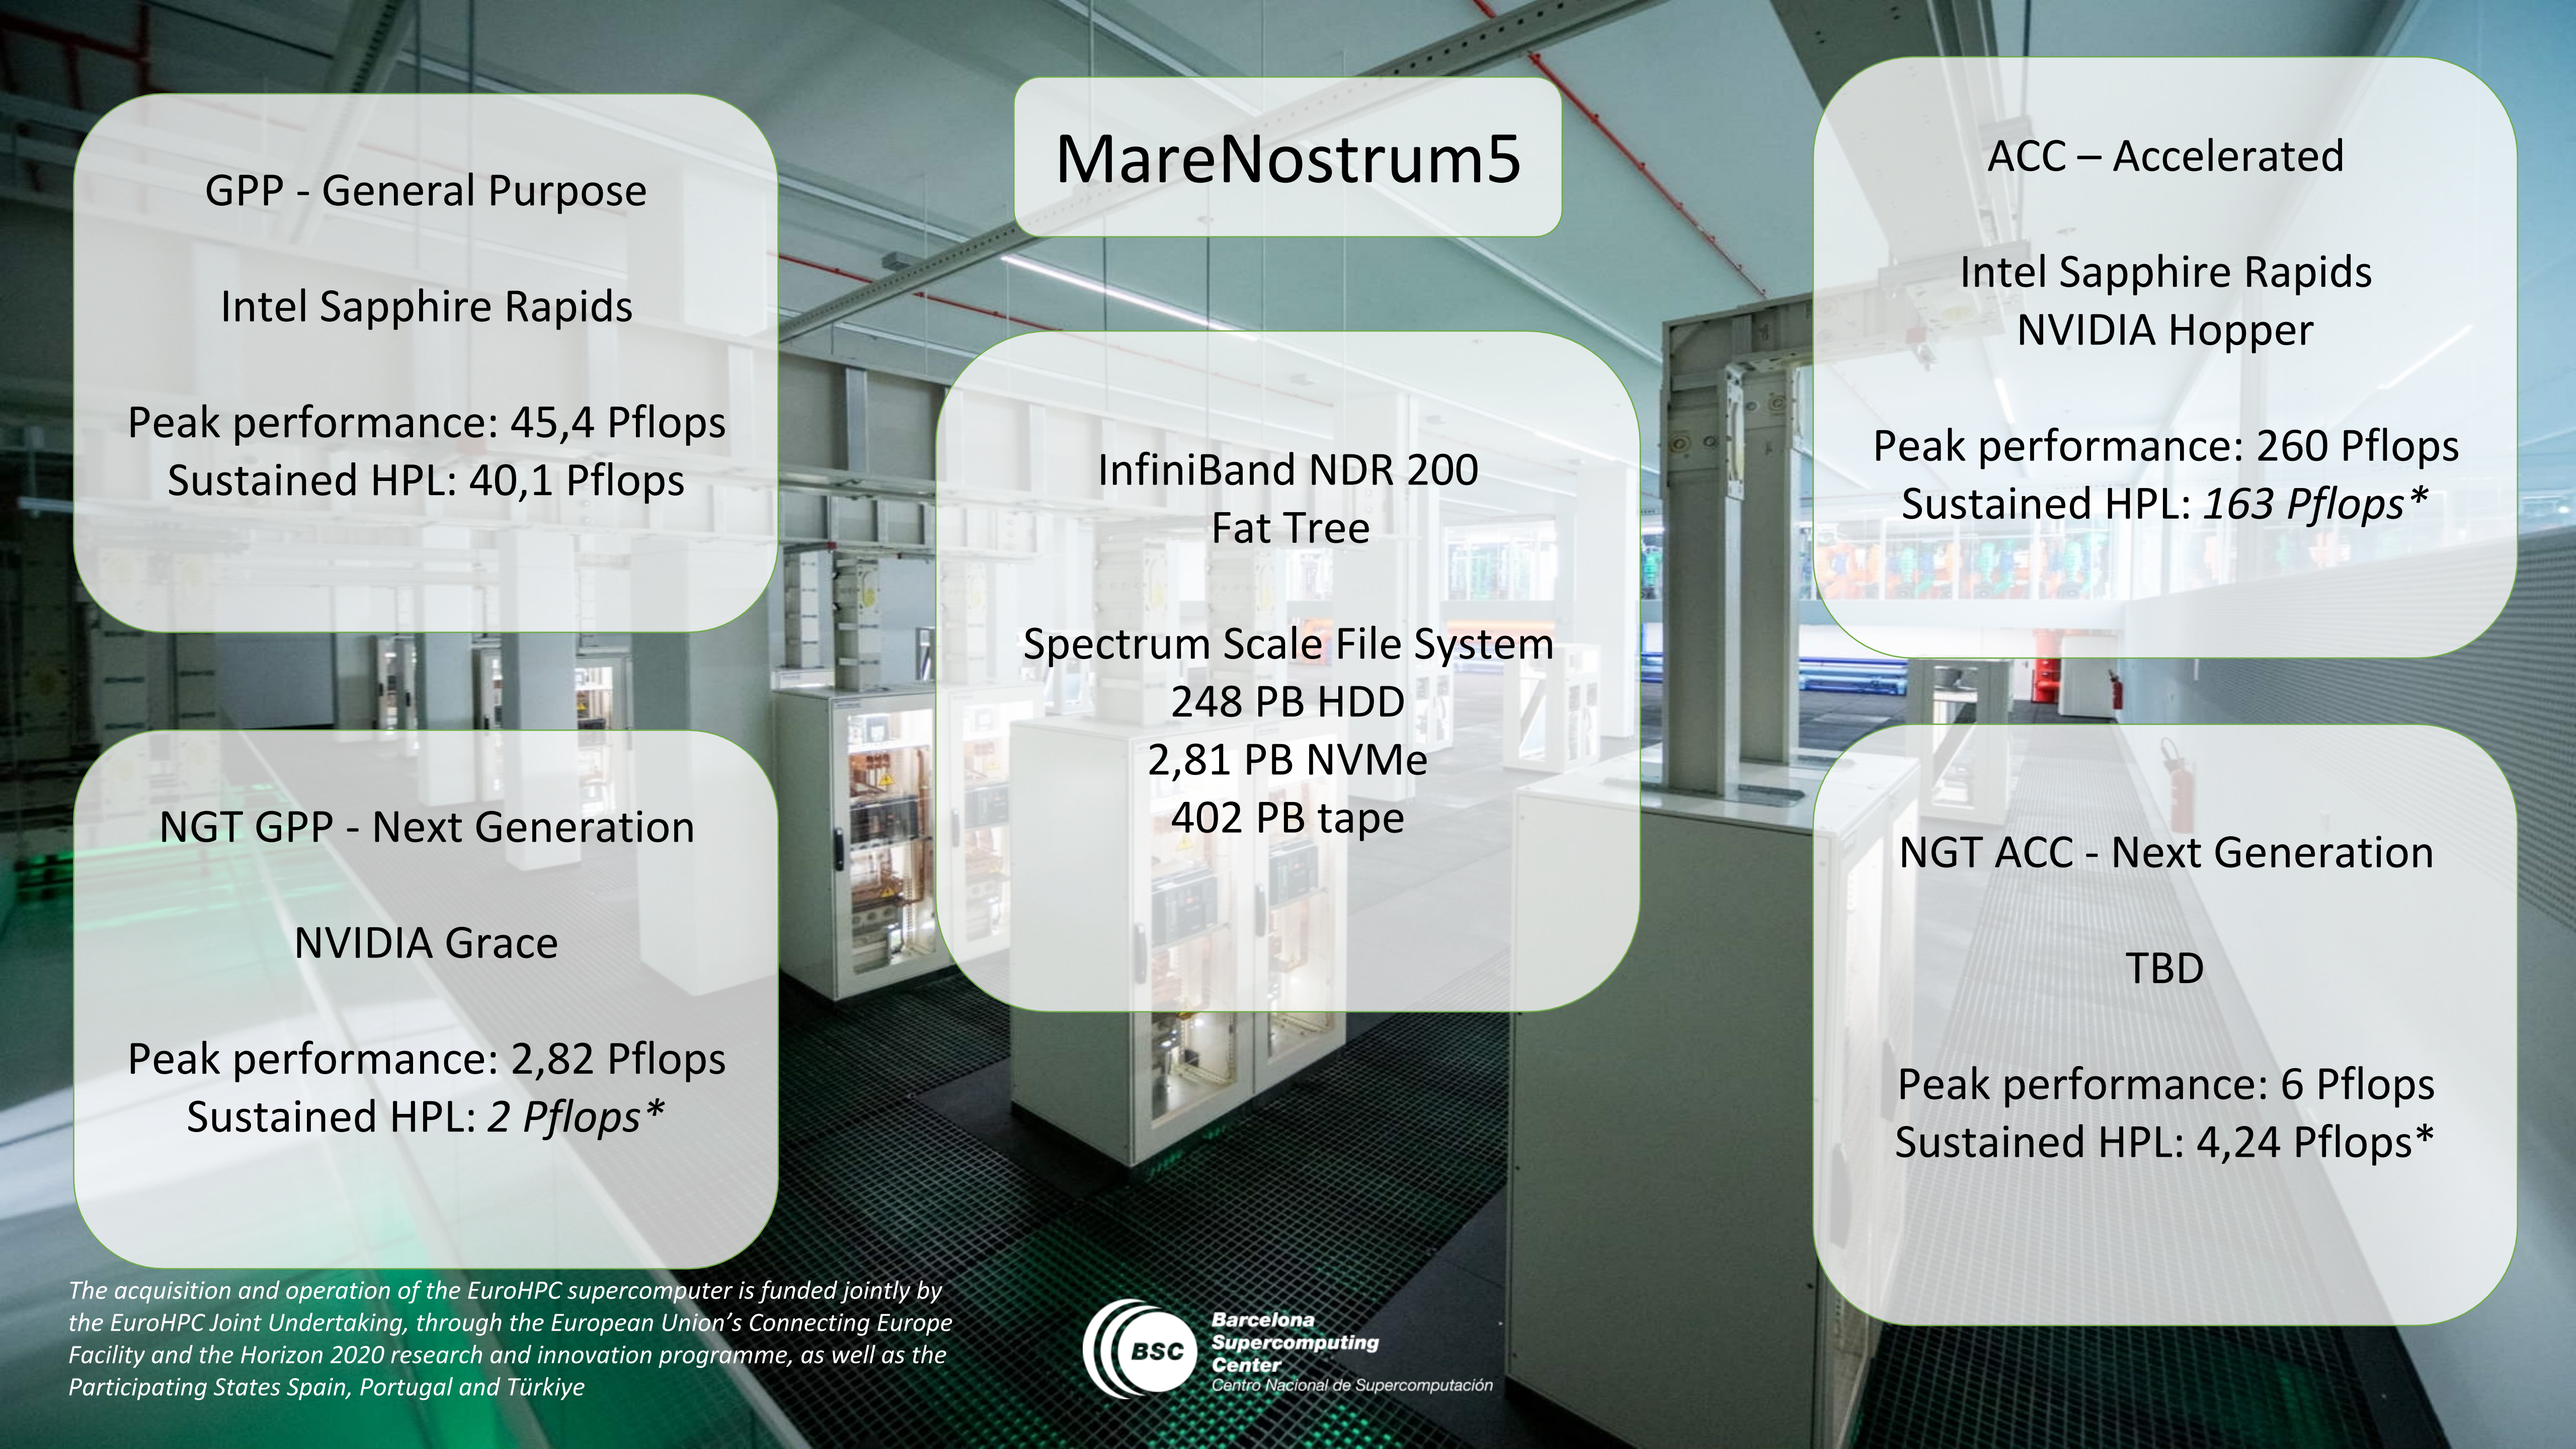

MareNostrum5

GPP - General Purpose

Intel Sapphire Rapids

Peak performance: 45,4 Pflops
Sustained HPL: 40,1 Pflops

InfiniBand NDR 200
Fat Tree

Spectrum Scale File System
248 PB HDD
2,81 PB NVMe
402 PB tape

ACC – Accelerated

Intel Sapphire Rapids
NVIDIA Hopper

Peak performance: 260 Pflops
Sustained HPL: *163 Pflops**

NGT GPP - Next Generation

NVIDIA Grace

Peak performance: 2,82 Pflops
Sustained HPL: *2 Pflops**

NGT ACC - Next Generation

TBD

Peak performance: 6 Pflops
Sustained HPL: 4,24 Pflops*

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

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

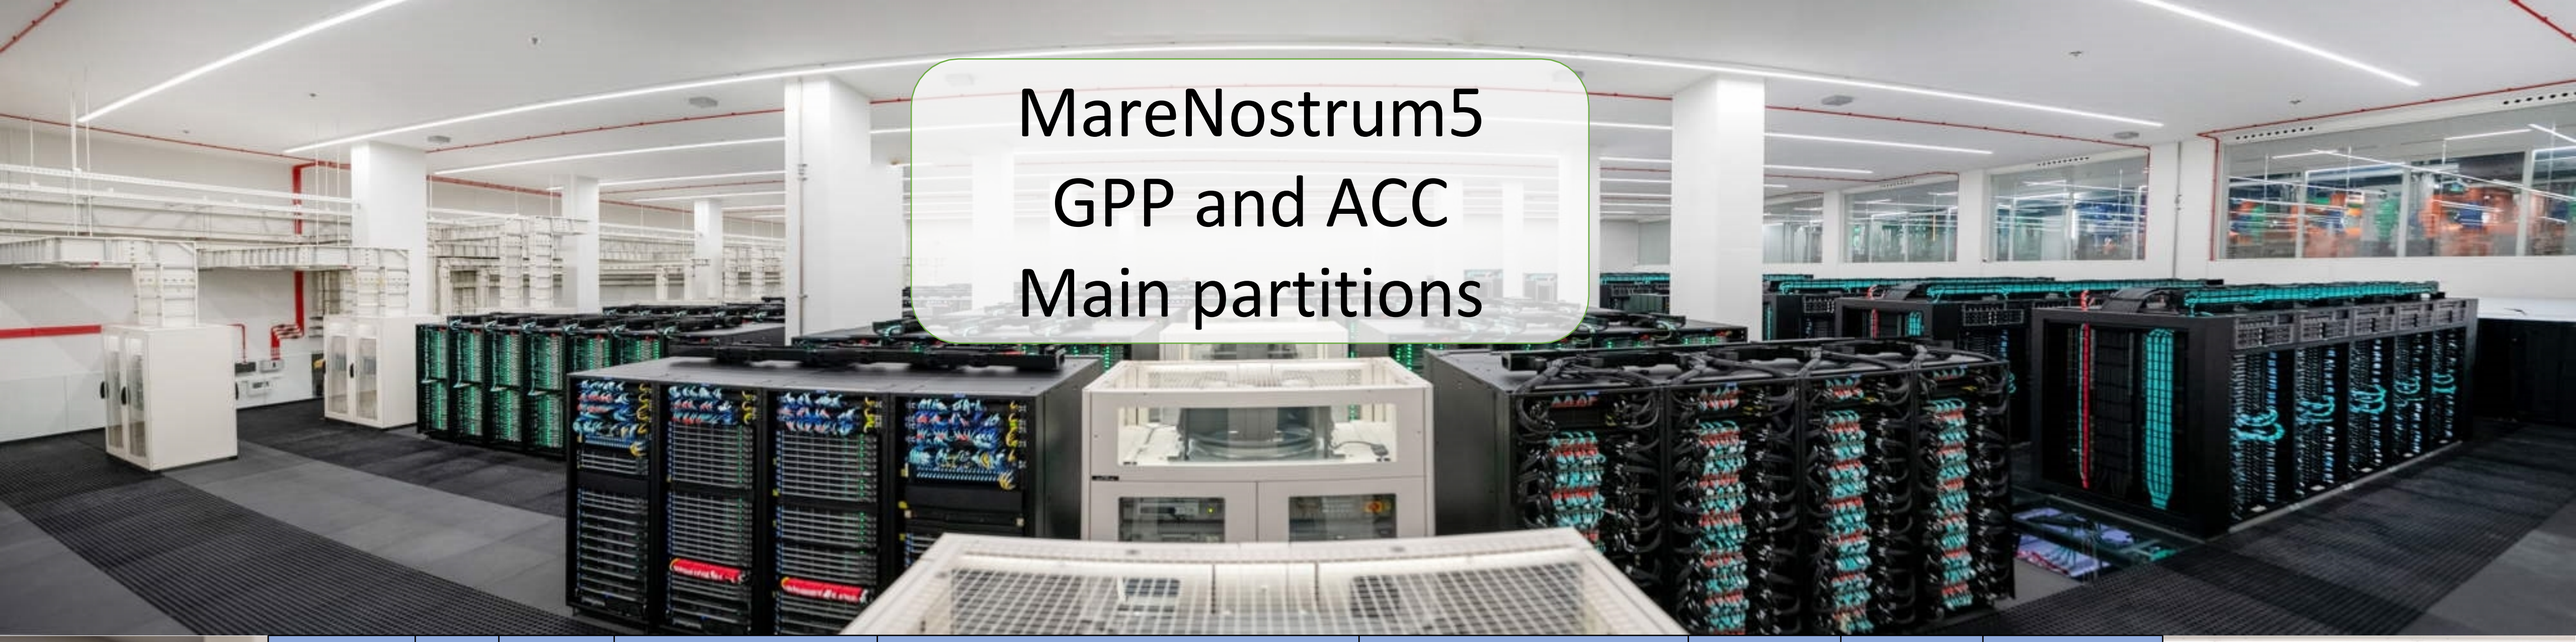

MareNostrum5 GPP and ACC Main partitions

Partition	Racks	Cooling	Nodes		Processor/Accelerator		Memory	PFlops (HPL)	Local Drive	High-Perf. Network
			Total	per rack						
GPP	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			
GPP HBM	1		72		2x Intel Sapphire R. 9480	56c @ 1.9GHz	> 0.5GB HBM/core 128GB HBM + 32GB DDR5	0.34		
ACC	35	DLC	1120	32	2x Intel Sapphire R. 8460Y+	40c @ 2GHz	512 GB DDR5	163* (138)	480GB NVMe	4 x NDR200
					4x Nvidia Hopper 64GB HBM2e					

MareNostrum 5

Total peak performance:	**314** Pflops
GPP:	45.4 Pflops
ACC:	260 Pflops
NGT GPP	2.82 Pflops
NGT ACC :	6 Pflops

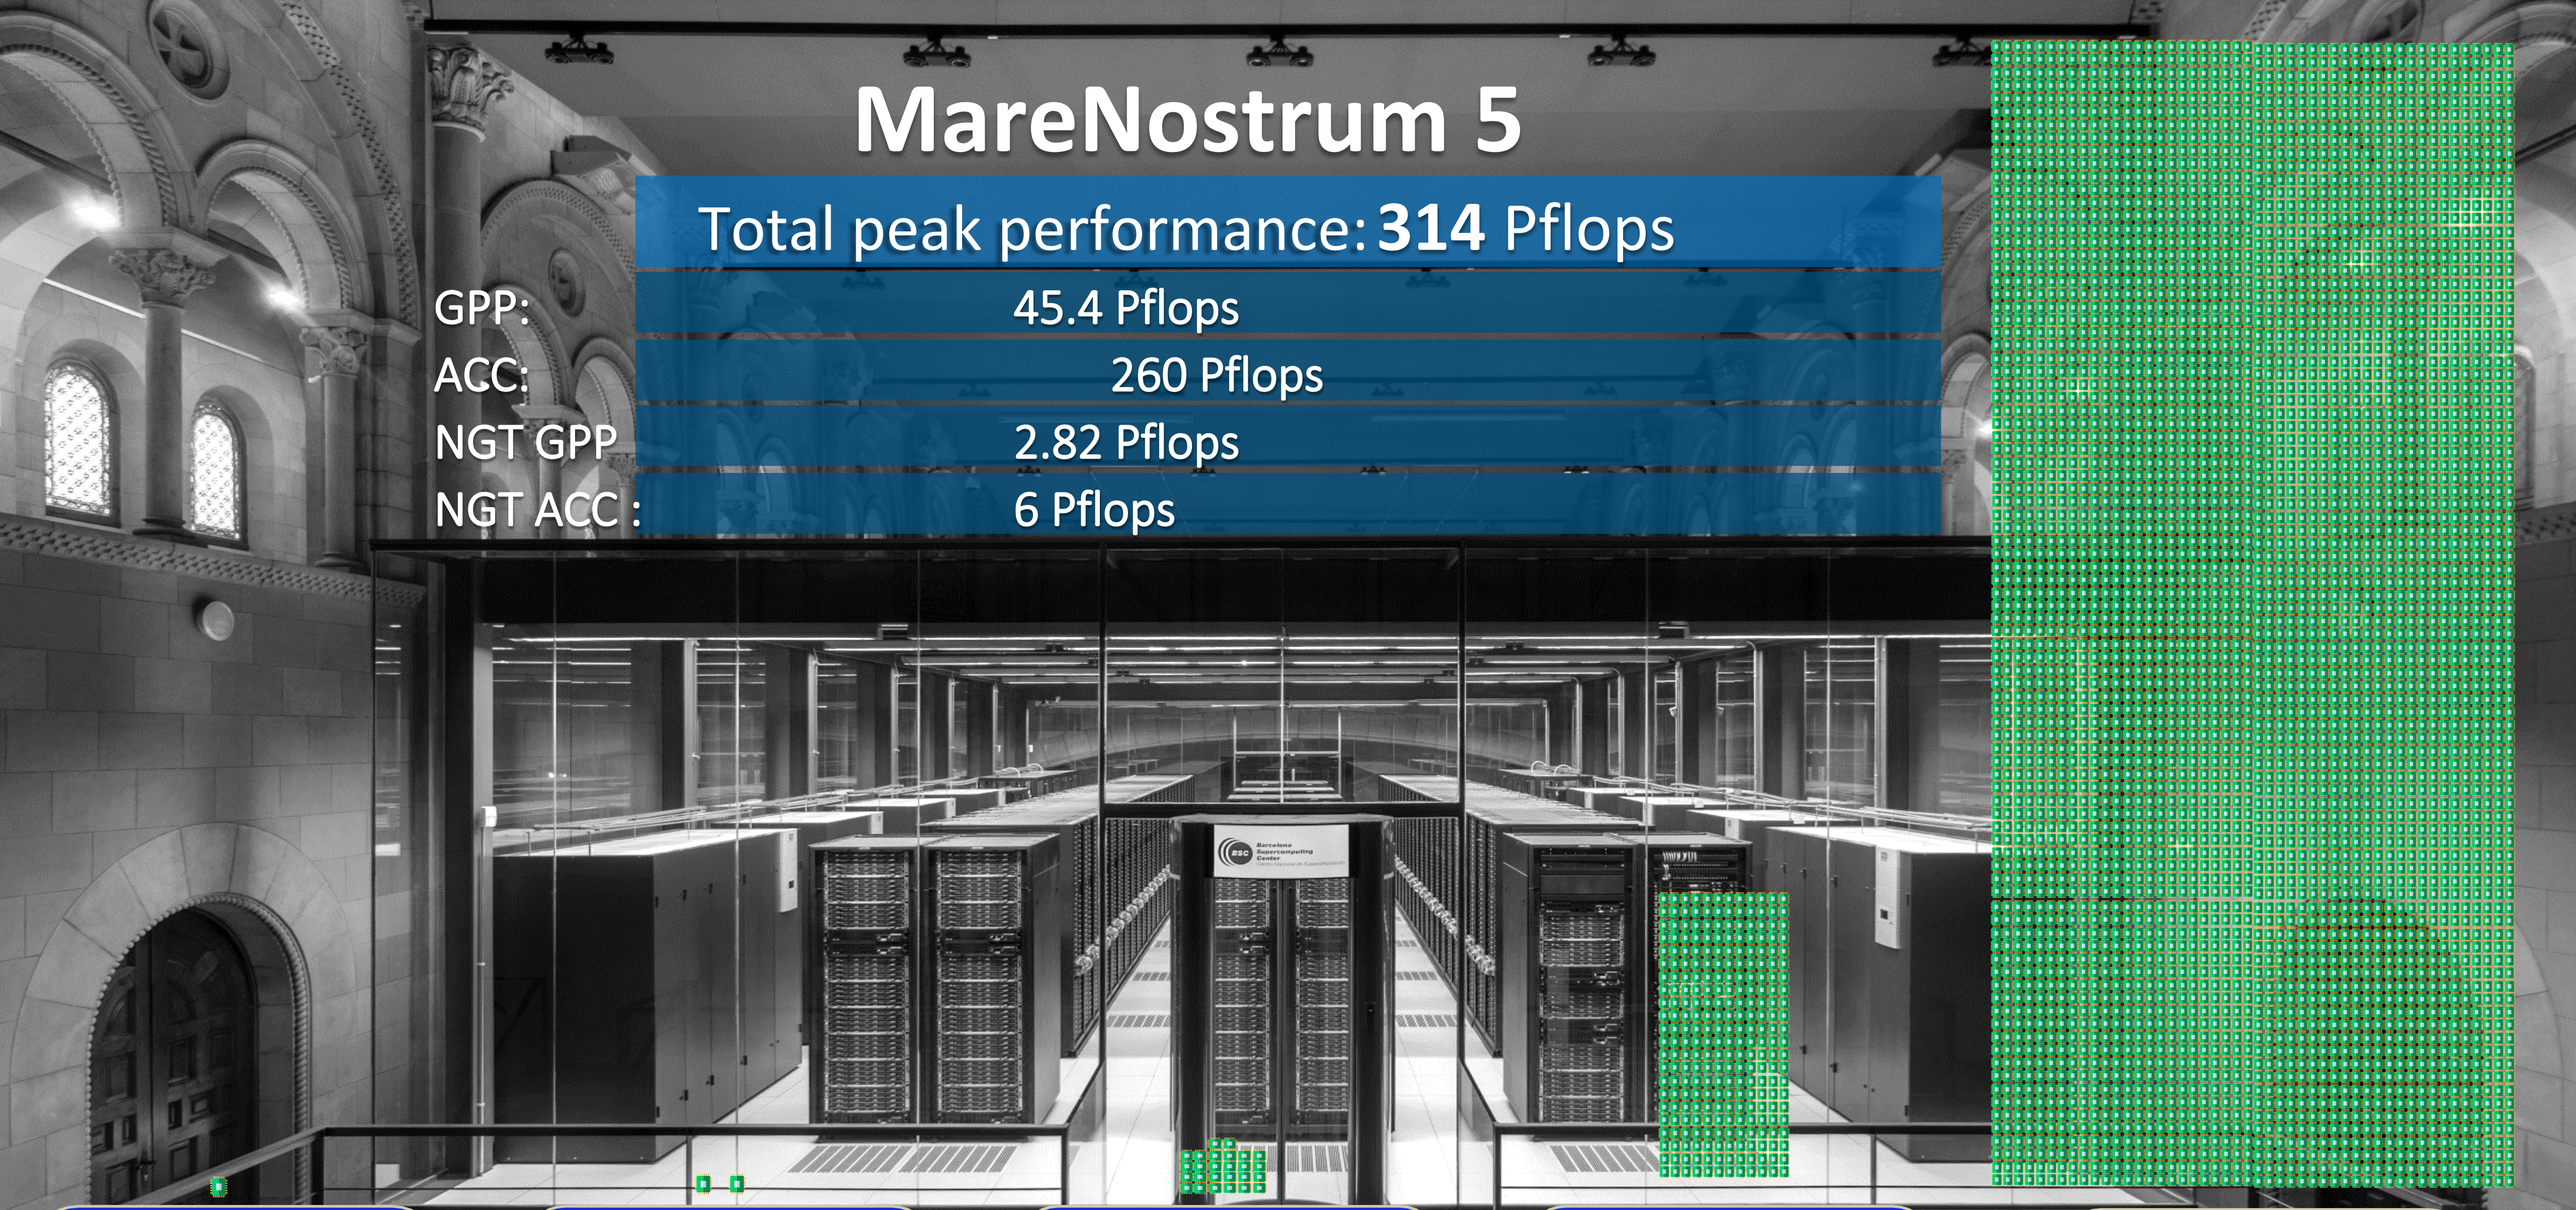

MareNostrum 1
2004 – 42.3 Tflops
1st Europe / 4th World
New technologies

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
2023 – 204.6 Pflops
3rd and 7th Europe
8th and 19th World

MareNostrum 4 – current usage

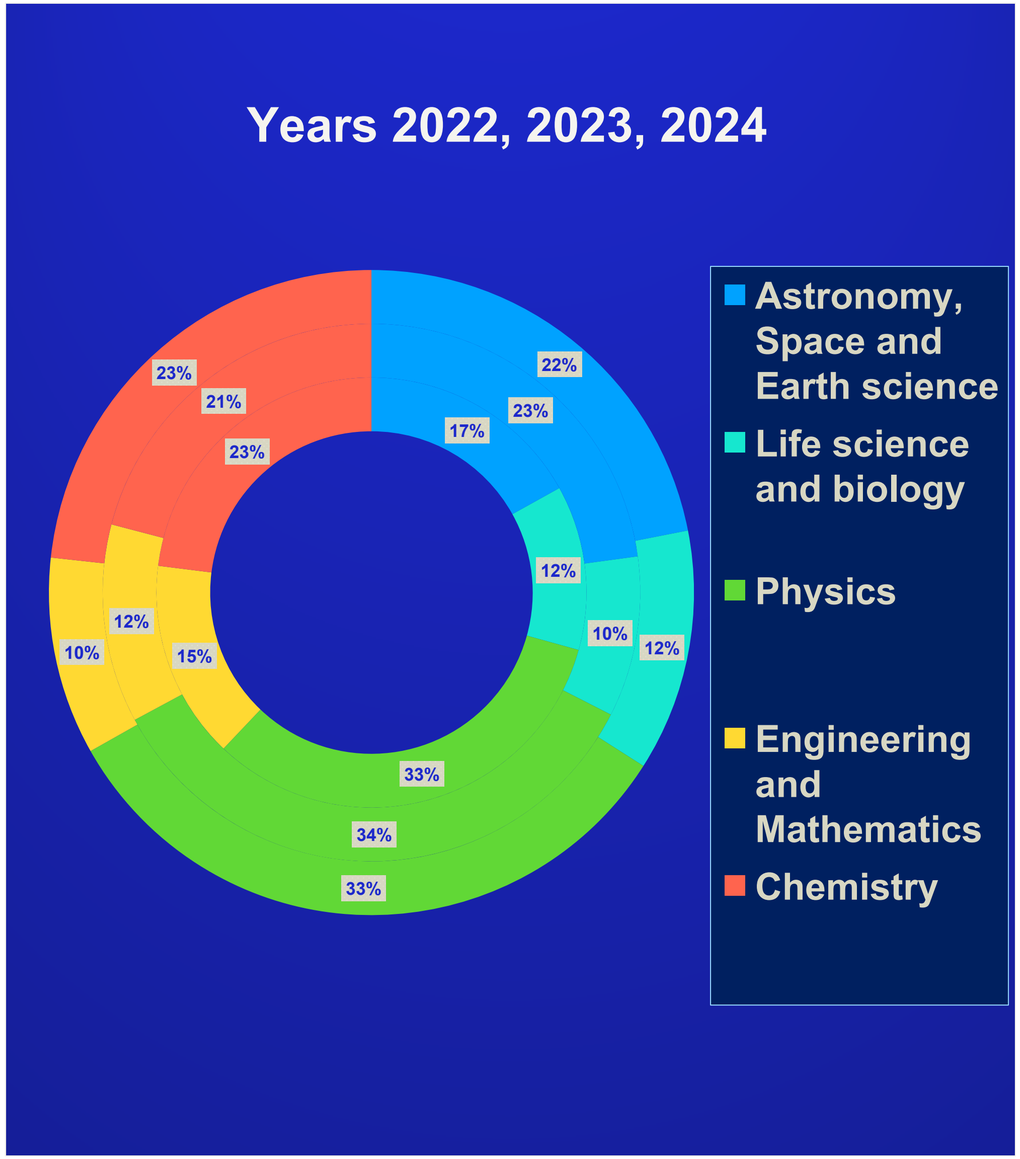

- Machine usage across different areas is fairly balanced, but depends on the annual calls.
- In the case of MN5, we anticipate a comparable distribution, with an increased utilization in specific domains including:
  - Artificial intelligence (ACC)
  - Large language models (ACC)
  - Large climate models (GCC + ACC)
  - Life science codes (GPP + ACC)
  - Large engineering runs to solve extremely large problems, like combustion (GPP)
  - Material Science (GPP + ACC)

MareNostrum 5 – current status

GPP and ACC partition expected pre-production on April 2024

EuroHPC Users expected to start using the machine on May 2024.

MN5 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.
- 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 and 400 PB on tapes, and more than 1.2 TB/s read-write bandwidth.

MareNostrum 5 – current status GPP

During the acceptance procedure for the GPP partition, the following observations were noted:

1. **Stability:** Great stability, during a 7-hour continuous run of the same job, over 99.999% job success with minimal variability.

2. **Scalability:** good scalability. Each island without contention (full fat tree) has 30 racks (2160 nodes), larger applications might face contention issues for high BW MPI calls.

3. **Storage Performance:** high performance for parallel fileystem. However, for specific applications, using local file system proved advantageous. The transfer of data from GPFS to the local filesystem should use Infiniband.

4. **Turbo Boost System:** The new turbo boost system from Intel Shappire Rapids processors performed very well, delivering very good performance/watt ratios.

```
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
APP: nemo
Config: Best energy
Median time: 1250.630 MN5-GPP
Maximum expected time +5%: 1252.0 -- 1314.6
Maximum expected energy +5%: 676.0 -- 709.8
 ./vendors-be-nemo-935027.err  MN5-GPP - Time: 1251.12  Energy: 637.86
 ./vendors-be-nemo-935029.err  MN5-GPP - Time: 1247.94  Energy: 636.23

--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
APP: nemo
Config: Best time
Median time: 306.376 MN5-GPP
Maximum expected time +5%: 328.0 -- 344.4
Maximum expected energy +5%: 709.0 -- 744.5
 ./vendors-bt-nemo-935026.err  MN5-GPP - Time: 307.081  Energy: 607.18
 ./vendors-bt-nemo-935028.err  MN5-GPP - Time: 306.328  Energy: 600.98

--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
APP: qe
Config: Best energy
Median time: 569.797 MN5-GPP
Maximum expected time +5%: 693.0 -- 727.6
Maximum expected energy +5%: 2734.0 -- 2870.7
 ./vendors-be-qe-935047.err  MN5-GPP - Time: 572.324  Energy: 1833.02
 ./vendors-be-qe-935049.err  MN5-GPP - Time: 569.553  Energy: 1789.13

--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
APP: qe
Config: Best time
Median time: 309.024 MN5-GPP
Maximum expected time +5%: 362.0 -- 380.1
Maximum expected energy +5%: 2856.0 -- 2998.8
 ./vendors-bt-qe-935046.err  MN5-GPP - Time: 309.028  Energy: 1914.27
 ./vendors-bt-qe-935048.err  MN5-GPP - Time: 309.274  Energy: 1922.58

--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
APP: wrf
Config: Best time
Median time: 350.033 MN5-GPP
Maximum expected time +5%: 396.0 -- 415.8
Maximum expected energy +5%: 405.0 -- 425.2
 ./vendors-bt-wrf-935066.err  MN5-GPP - Time: 350.442  Energy: 345.111
 ./vendors-bt-wrf-935080.err  MN5-GPP - Time: 351.494  Energy: 341.048

--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
APP: wrf
Config: Best energy
Median time: 1188.750 MN5-GPP
Maximum expected time +5%: 1387.0 -- 1456.4
Maximum expected energy +5%: 354.0 -- 371.7
 ./vendors-be-wrf-935067.err  MN5-GPP - Time: 1186.98  Energy: 304.766
 ./vendors-be-wrf-935069.err  MN5-GPP - Time: 1188.75  Energy: 309.994

--------------------------------------------------------------------------------
```

EUROHPC SUMMIT 2024

ANTWERP

Thanks !

David Vicente
BSC – CNS
david.vicente@bsc.es

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

UNLEASHING THE POWER OF EUROPEAN HPC AND QUANTUM COMPUTING

EUROHPC SUMMIT 2024

ANTWERP

UNLEASHING THE POWER OF EUROPEAN HPC AND QUANTUM COMPUTING

EuroHPC Systems: Current Usage, Anticipated Needs and Future Challenges

Leonardo, Cineca

Daniela Galetti

Leonardo partitions

- Now 6th on Top500: 241 PF + 8 PF
- TCO Investment: 240M€
- Direct Liquid Cooled: 95%
- Warm water: 37- 47 Celsius degrees
- Data space: more than 100PB (NVMe+HDD)
- Interconnect: IB HDR 200
- Topology: Dragonfly+

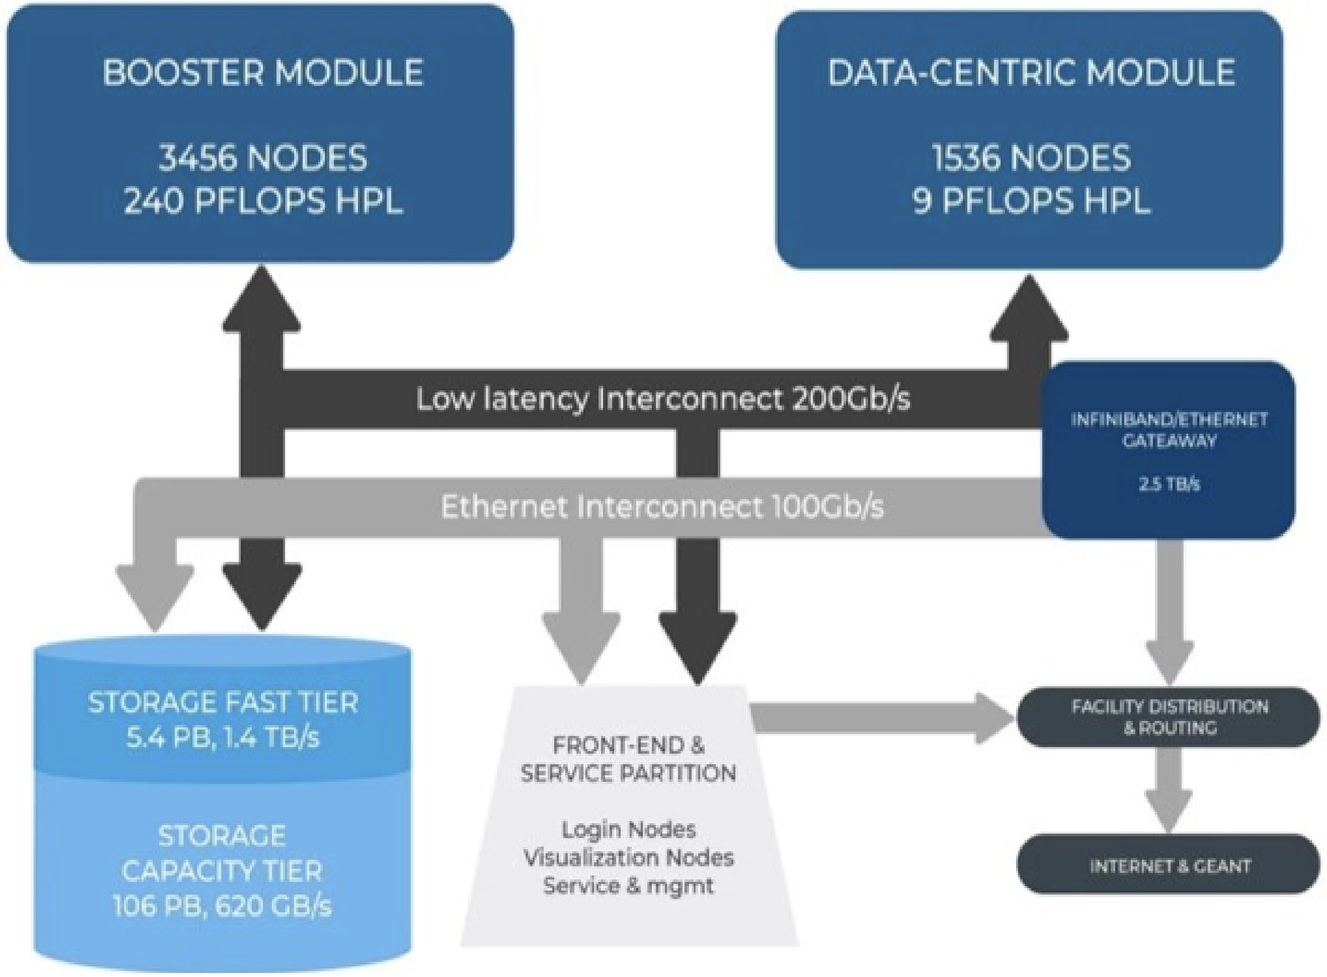

Current Usage: Leonardo access and projects

- 2730 Active Users

- 1124 Defined Projects
  - 866 Active
  - 250 Expired (some special expired projects can still run, but in low priority)
  - 2 Closed
  - 6 Removed

Current Usage: Leonardo workloads

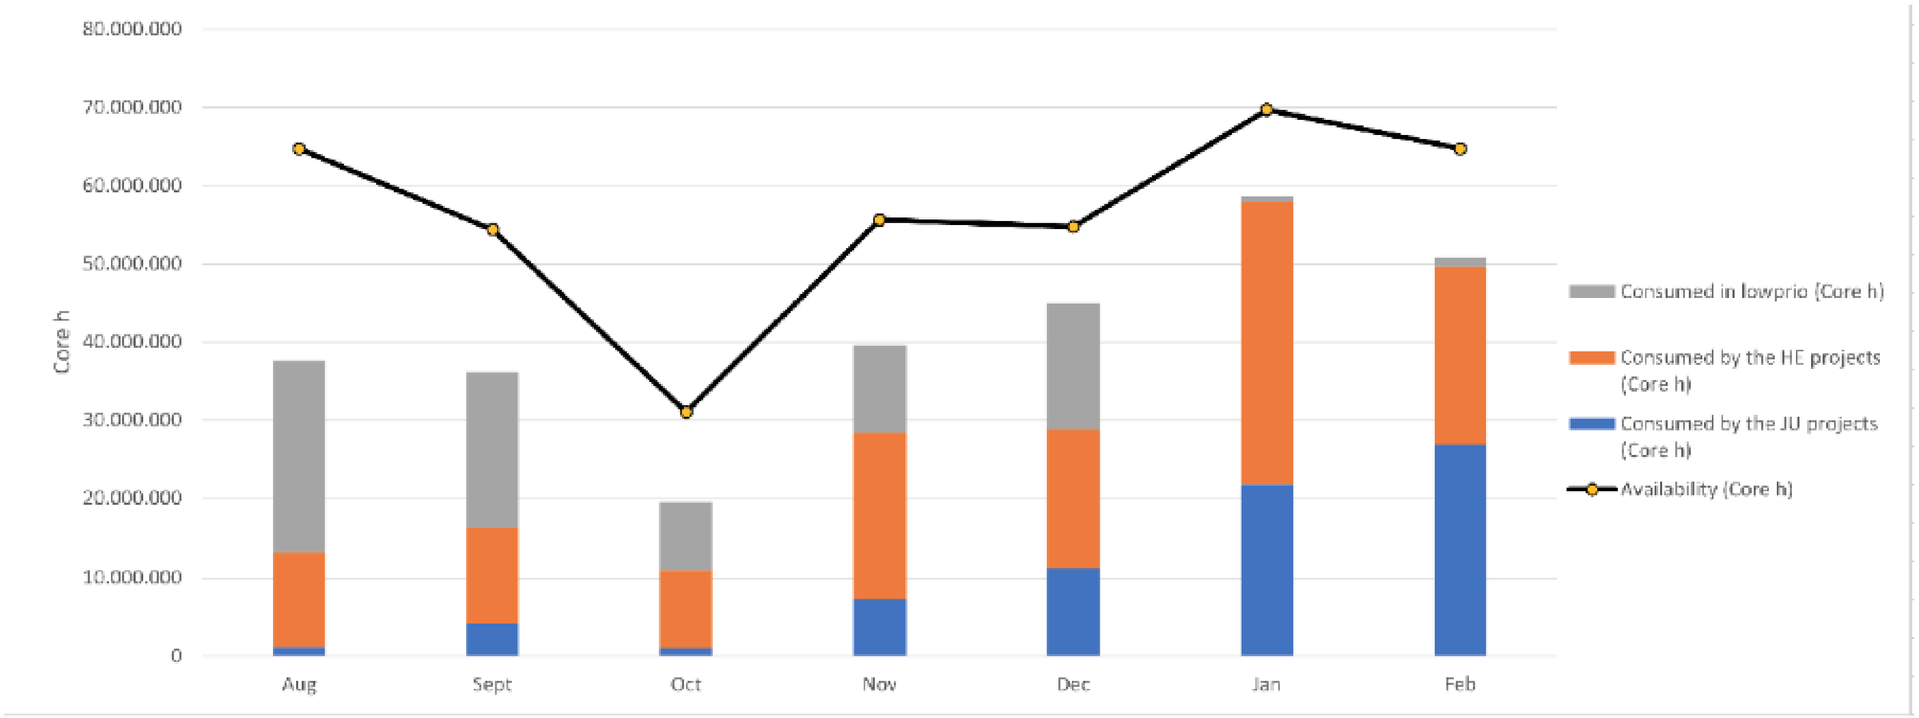

Leonardo Booster Partition Usage

Type of applications by cpu hours budget

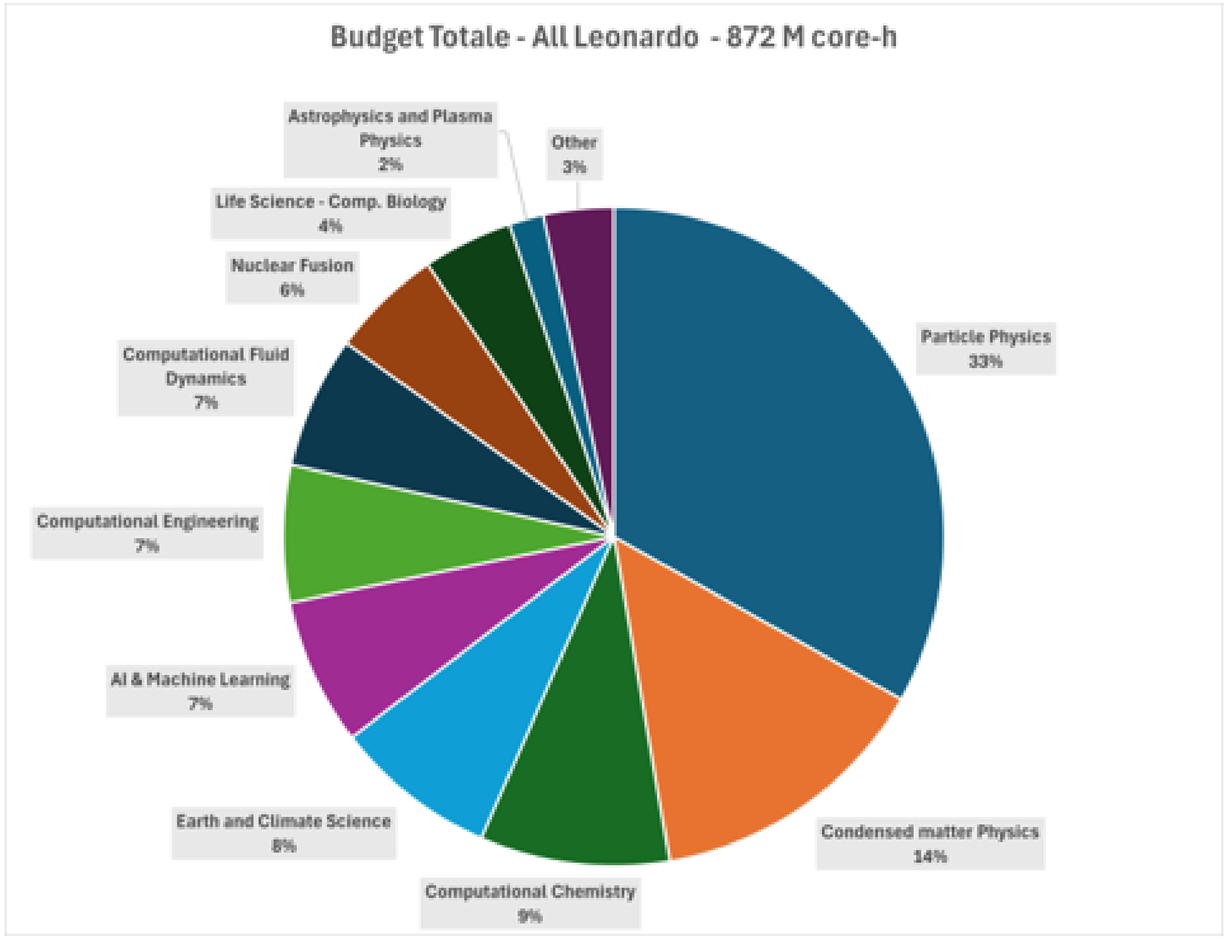

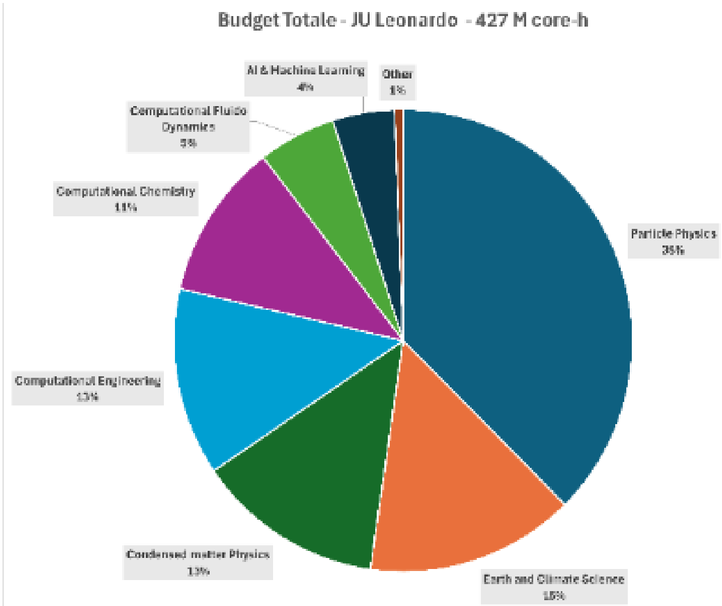

Leonardo Booster Partition Usage

Type of applications by amount of projects

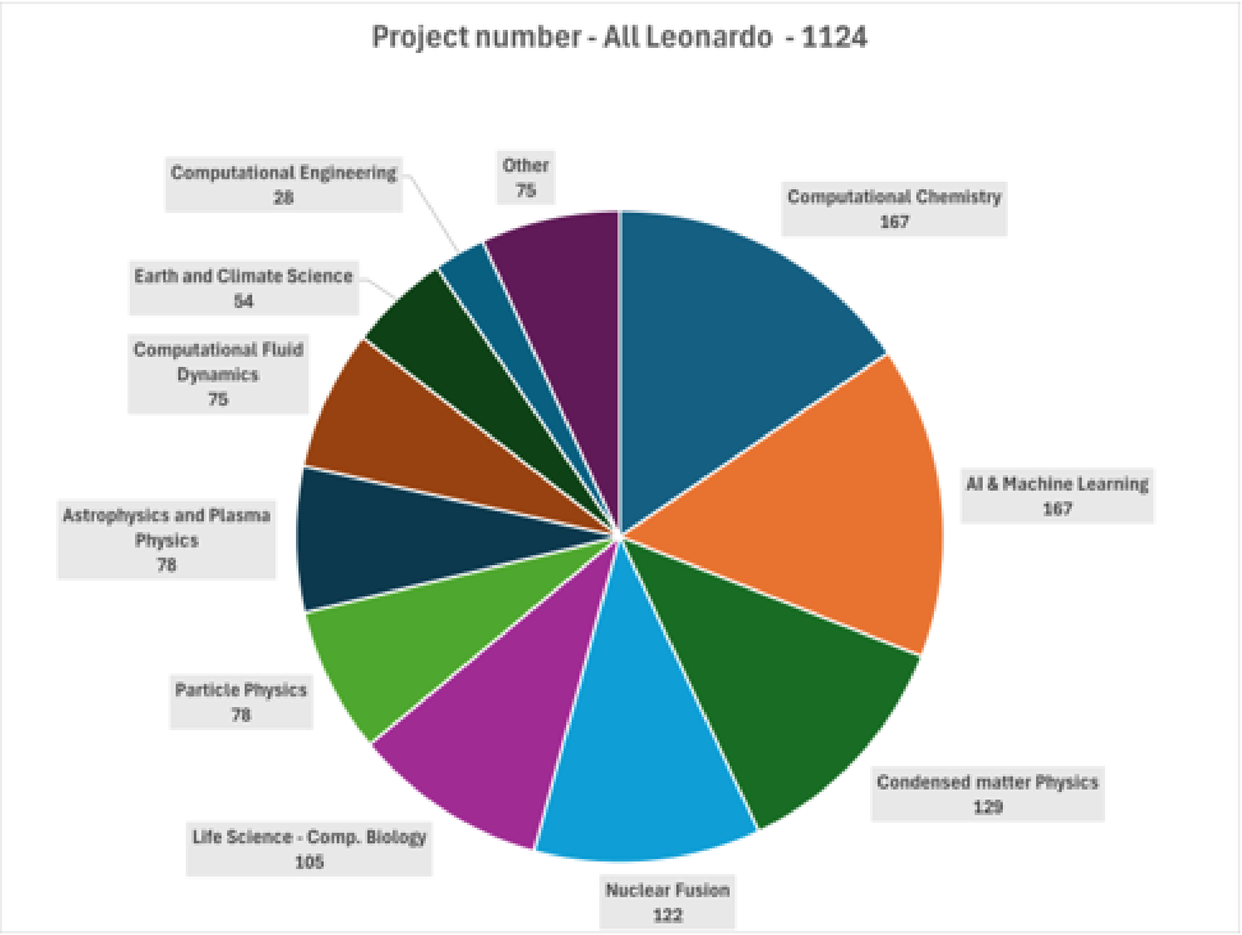

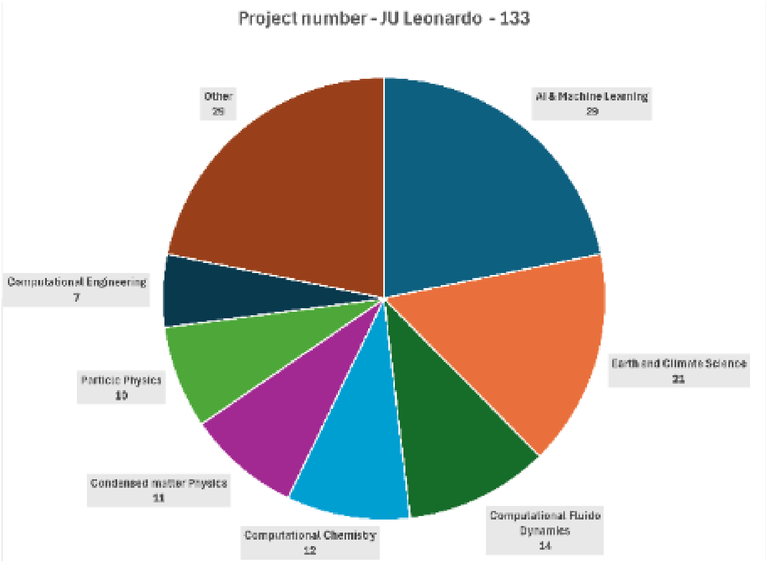

Anticipated needs and future challenges

- Data Centric General Purpose (DCGP) partition production consolidation
- Generative AI
- Fast storage
- Capacity computing and/or urgent computing
- Quantum computing

EUROHPC SUMMIT 2024

ANTWERP
18-21 MARCH

TO EXASCALE AND BEYOND

Thank You!

EUROHPC SUMMIT 2024

ANTWERP

EuroHPC Systems: JUPITER

Joint **U**ndertaking **P**ioneer for **I**nnovative and **T**ransformative **E**xascale **R**esearch

UNLEASHING THE POWER OF EUROPEAN HPC AND QUANTUM COMPUTING

Benedikt von St. Vieth (b.von.st.vieth@fz-juelich.de, jupiter@fz-juelich.de), Jülich Supercomputing Centre

JUPITER System A

- Modular System Architecture
- JUPITER Management Stack
- Eviden SMC xScale
- ParaStation Modulo
- Slurm (Scheduler)
- EasyBuild (Scientific Software)

- >1 ExaFLOP/s FP64 HPL
- >70 ExaFLOP/s FP8 AI
- ~6000 Nodes
- ~24000 NVIDIA Grace-Hopper
- NVIDIA Mellanox NDR
- BullSequana XH3000

- >5 PetaFLOP/s FP64 HPL
- ~1300 Nodes
- ~2600 SiPearl Rhea1
- BullSequana XH3000

- 210 PB useable SAS
- ~1,5 TB/s
- 22* IBM SSS6000

- 23 Login Nodes
- NVIDIA Hopper/A40
- 200 Gbit/s connectivity

- 21 PB useable Flash
- ~2,5 TB/s
- 20* IBM SSS6000

Core Configuration

Future Call Modules

Research Fields in Operation

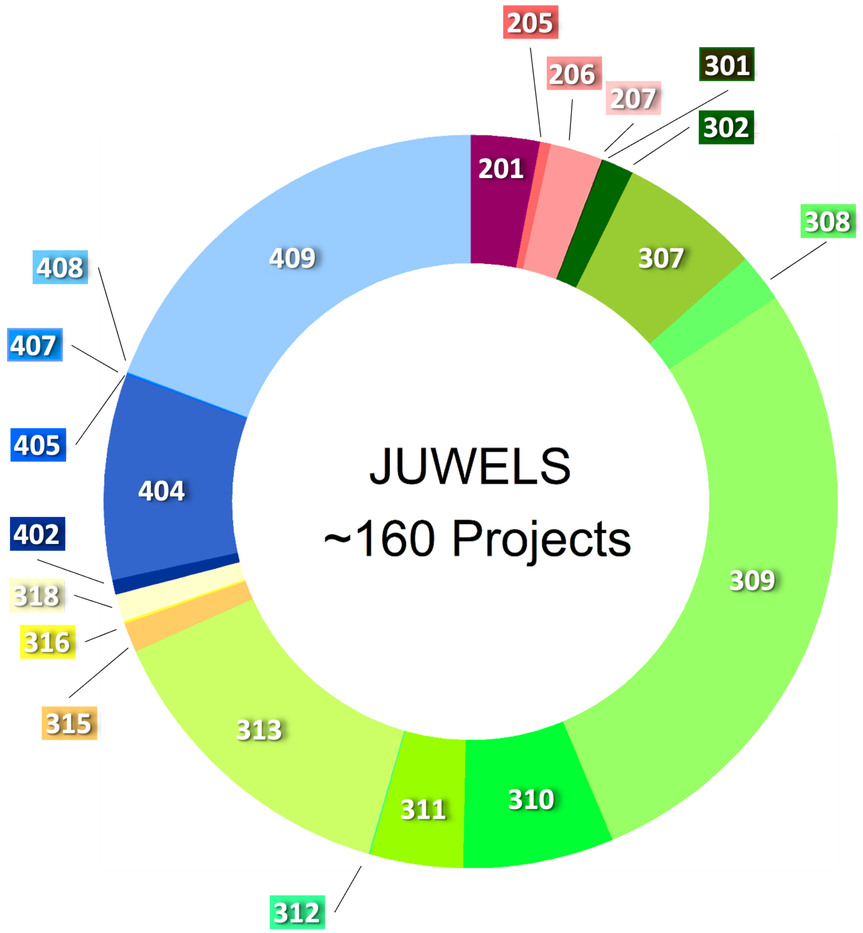

Reseach Fields

- 201 Basic Biological and Medical Research
- 205 Medicine
- 206 Neurosciences
- 207 Agriculture, Forestry and Veterinary Medicine
- 301 Molecular Chemistry
- 302 Chemical Solid State and Surface Research
- 307 Condensed Matter Physics
- 308 Optics, Quantum Optics and Physics of Atoms, Molecules and Plasmas
- 309 Particles, Nuclei and Fields
- 310 Statistical Physics, Soft Matter, Biological Physics, Nonlinear Dynamics
- 311 Astrophysics and Astronomy
- 312 Mathematics
- 313 Atmospheric Science, Oceanography and Climate Research
- 315 Geophysics and Geodesy
- 316 Geochemistry, Mineralogy and Crystallography
- 318 Water Research
- 402 Mechanics and Constructive Mechanical Engineering
- 404 Heat Energy Technology, Thermal Machines, Fluid Mechanics
- 405 Materials Engineering
- 407 Systems Engineering
- 408 Electrical Engineering and Information Technology
- 409 Computer Science

This is the current status on JUWELS.

We expect similar distribution at least on the national part of JUPITER.

Research Fields at the Procurement

		Booster			Cluster	MSA
Name	Field	GPU	GPU High-Scale	CPU	CPU	
Amber	Molecular Dynamics	✓				
Arbor	Neuroscience	✓	✓			
Chroma	QCD	✓	✓			
Gromacs	Molecular Dynamics	✓				
ICON	Climate	✓				
JUQCS	QC	✓	✓			✓
nekRS	CFD	✓	✓			
ParFlow	Earth System	✓				
PIConGPU	Plasma	✓	✓			
Quantum ESPRESSO	Material Science	✓				
SOMA	Polymer Systems	✓				
AI-MMoCLIP	AI (Mixed mode)	✓				
AI-NLP	AI (LLM)	✓				
AI-ResNet	Computer Vision	✓				
dynQCD	QCD				✓	
NAStJA	Biological				✓	
Graph500	Graph Traversal			✓		
HPCG	Conjugate Gradient	✓			✓	
HPL	Linear Equations	✓			✓	
IOR	Filesystems/IO			✓	✓	
LinkTest	Network			✓	✓	✓
OSU	Network	✓		✓	✓	
STREAM	Memory	✓			✓	

JUPITER Research and Early Access Program

IMPORTANT: Timeline not final yet!

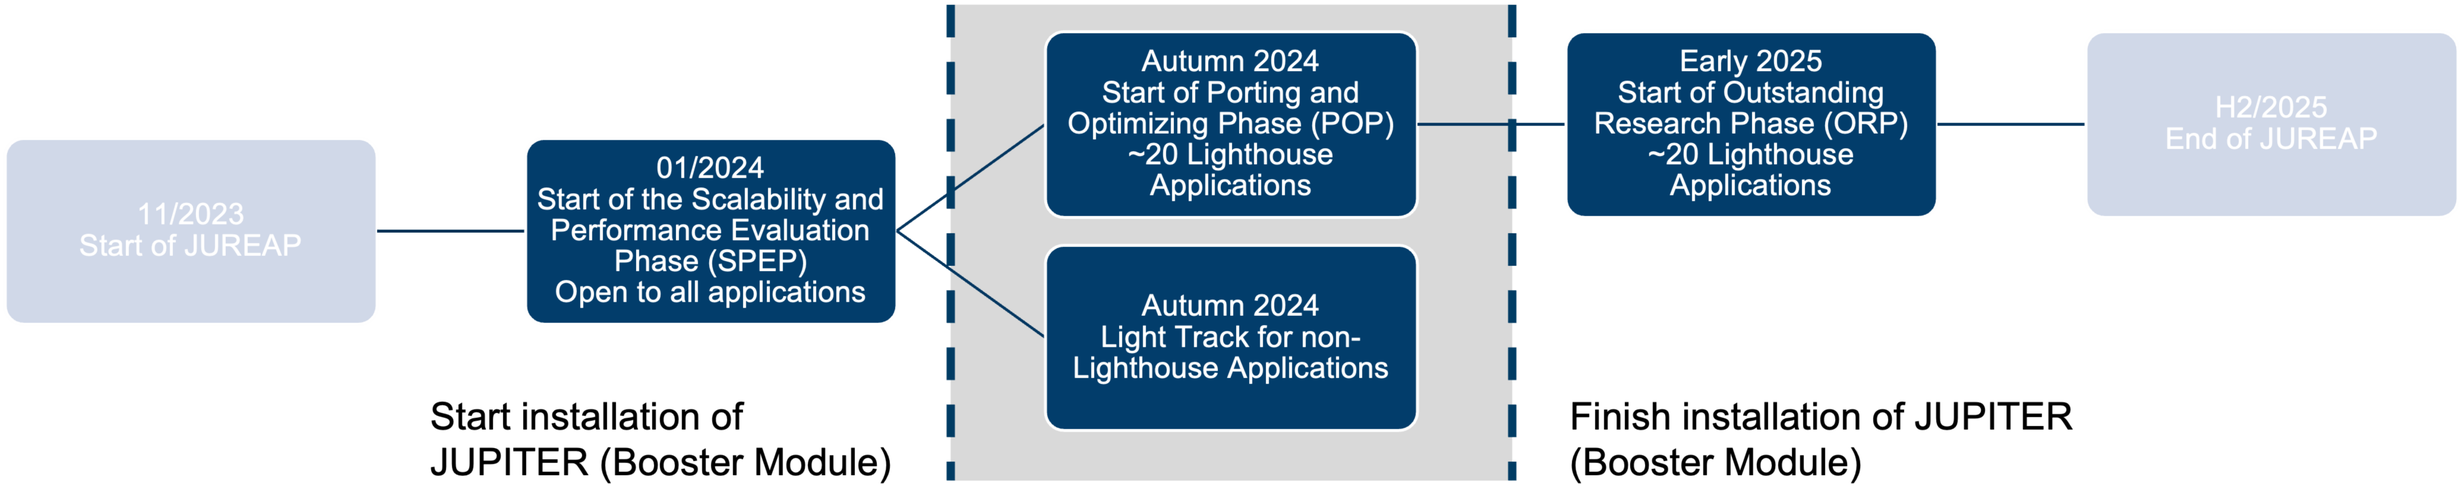

Phase 1: Scalability and Performance Evaluation Phase (SPEP)
Phase 2: Porting and Optimizing Phase (POP)
Phase 3: Outstanding Research Phase (ORP)

JUREAP @ #EXA_jupiter

Seeding Exascale in Europe!

Contact: jureap@fz-juelich.de
Mailing list: https://lists.fz-juelich.de/mailman/listinfo/jureap-info

JÜLICH Forschungszentrum

EUROHPC SUMMIT 2024

ANTWERP

UNLEASHING THE POWER OF EUROPEAN HPC AND QUANTUM COMPUTING

EuroHPC systems : current usage, anticipated needs and future challenges

Views on the Jules Verne side

S. Requena (GENCI)

GENCI, a French HPC research infrastructure

- Serving yearly 1700 research projects in HPC and AI (academia, industry)

TGCC/CEA - Ile de France

- **Hosting Site** for the 2nd Exascale system (**EuroHPC**) with Jules Verne
- Hosting Site for the **1st hybrid HPC + Quantum computing infrastructure** (HQI, HPCQS, EuroQCS-France)

IDRIS/CNRS - Ile de France

- **1st converged HPC/AI system** *(#AIForHumanity)*
- **Sovereign** computing facilities / services for the AI community
- **>900 yearly projects in AI**
- **> 3100 GPUs and more to come soon** 🚀

CINES/FU - Montpellier

- **-> 70 PF** with AMD next gen GPUs (>1400) & CPUs (>100k)
- **Next step before French Exascale system**

#3 THE GREEN 500

EUROHPC SUMMIT 2024

ANTWERP 18-21 MARCH

TO EXASCALE AND BEYOND

Exascale is the NEW frontier

SKAO

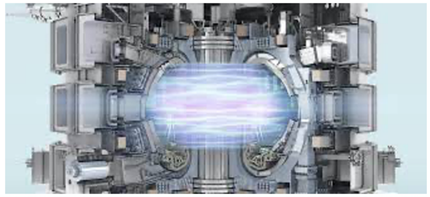

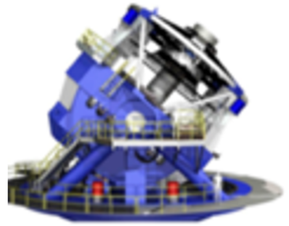

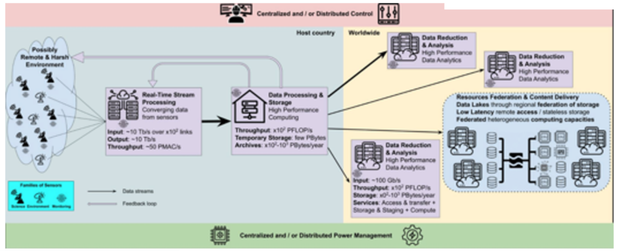

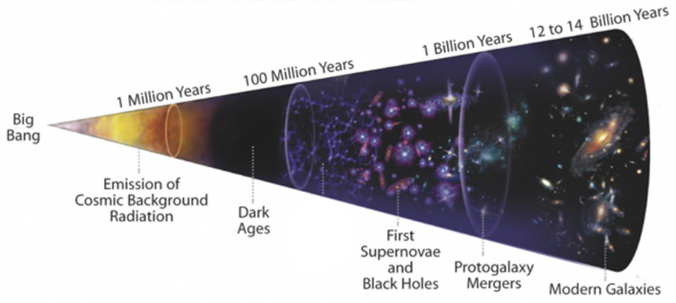

SAFRAN

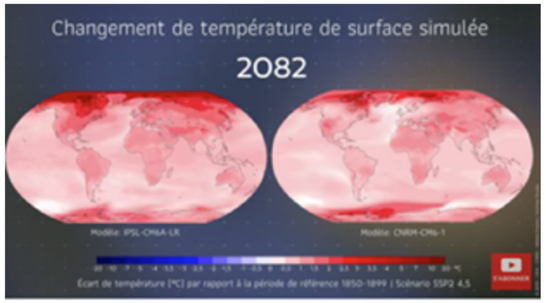

EuroHPC Joint Undertaking

beEU belgium24.eu

Jules Verne : The French led Exascale project

Organization of the french application

- GENCI *Hosting Entity*
- CEA *Hosting Site*
- SURF (NL) as member of consortium

Name of the consortium : Jules Verne

Full TCO over 5 years : 542 M€ (50% EuroHPC, 50% consortium)

- French public contribution
- NL contribution

-Seeking more partners on the consortium to reach 300M€

- **International partners**
- French research institutions
- French industrial partners (end users)

-System installed at TGCC (CEA) starting end 2025

EUROHPC SUMMIT 2024

ANTWERP 18-21 MARCH

TO EXASCALE AND BEYOND

Our vision for an Exascale system

Addressing societal and scientific challenges (such as universe sciences, climate change, health, new energy, innovative materials, transport or smart cities/systems) via large scale numerical simulations and massive data analysis using artificial intelligence

- An accelerator of European Science and Innovation
  open to all scientific and industrial collaborations, supporting new services including Cloud based interactive supercomputing / visualisation, containerisation and urgent computing for fast decision making (public, industrial)
- A converged HPC/HPDA/AI system with a modular, balanced and energy efficient architecture
  based on accelerated, scalar and HPDA partitions within a tiered data centric infrastructure
  integrating state-of-the-art post-exascale quantum accelerators and related services for specific workloads
- A system fully embedded inside the digital continuum
  ready for secured end-to-end workflows from instruments / edge devices to long term sovereign storage
- A system with European Technology and Skills
  integrating European hardware / software technologies in terms of computing, storage, network, infrastructure, middleware, applications with global support of AST to engage/support communities.

A system ready to harness European technologies and the best breed of opensource software in a highly secure environment

EuroHPC Joint Undertaking

beEU belgium24.eu

Final words : the challenges

- Porting HPC (legacy) applications to accelerated architectures
- Development of new services towards AI, scientific instruments, public services, Cloud-based access… including public-private continuum
- Integration of new (EU) architectures (concepts) : tight integration CPU/GPU – APU, aggregation / disaggregation of resources, coupling of QC accelerators and hybrid HPC/QC software stack…
- Energy efficiency / management (from components to applications)
- And last but not least : pan European training and user support to the (new) communities → **the sinews of war**